Envelope Surfaces, Surface Design and Meshing

Kruithof, Nico Gerard Hugo; Kruithof, N.GH

Envelope Surfaces

Surface Design and Meshing

Nico Gerard Hugo Kruithof

The research presented in this thesis was funded by the IST Programme of the European Community 'ECG – Effective Computational Geometry for Curves and Surfaces' (EU Project No.: IST-2000-26473) and 'ACS – Algorithms for Complex Shapes' (EU Project No.: IST-006413).

Front cover: Uitsnede uit het skin surface model van een caffeine molecule.
Ontwerp: Reinier van den Berg en Nico Kruithof

Rijks*universiteit* Groningen

Envelope Surfaces

Surface Design and Meshing

Proefschrift

ter verkrijging van het doctoraat in de
Wiskunde en Natuurwetenschappen
aan de Rijks*universiteit* Groningen
op gezag van de
Rector Magnificus, dr. F. Zwarts,
in het openbaar te verdedigen op
maandag 13 maart 2006
om 16.15 uur

door

Nico Gerard Hugo Kruithof

geboren op 17 juli 1978
te Dordrecht

Promotor:	Prof. dr. G. Vegter
Beoordelingscommissie:	Prof. dr. J-D. Boissonnat Prof. dr. T.K. Dey Prof. dr. M.H. Overmars

ISBN: 90–367–2516–X

Contents

Chapter 1

Introduction

Surfaces are important in research disciplines like Computer Aided Design, Computational Geometry and Computer Graphics.

Computer Aided Design [8, 52] is the development of a wide range of computer-based tools assisting in the design and creation of products and goods. These tools help, amongst others, engineers and architects in their design activities. In Computer Aided Design, the focus lies on construction and design of curves and surfaces.

Computational Geometry [60, 89] is the study of (computer) algorithms to solve problems stated in terms of geometry. Computational geometry has recently expanded its scope to include curves and surfaces. The research presented in this thesis is funded by EU projects dedicated to curves and surfaces and focusses on retrieval of properties of the curves and surfaces.

Computer Graphics [63] is the research field dedicated to visualization, where one utilizes computers both to generate visual images and to alter visual and spatial information sampled from the real world. Emphasis lies on visualization and output of surfaces.

Irrespective of the purpose for which the surface is used, without a suitable description a computer cannot be of any use. A *surface representation* is the formulation of a surface such that a human is able to reason about the surface with the aid of a computer. Many representations have been proposed [18, 49, 52] each with different advantages and disadvantages. Conversions between representations are needed in order to be able to use the advantages of different representations in the same application. E.g., piecewise linear approximations are regularly used for visualization and further geometric processing. Often it is not possible to give a representation that defines the same surface as the original surface. In such cases the surface is *approximated*.

Starting point of this thesis is the class of skin surfaces used in molecular biology [48]. Several methods exist that deform one skin surface efficiently into another skin surface [28, 33, 48] making them potentionally suitable for animations. We show how an arbitrary surface can be approximated with a skin surface in Chapter 3 and construct a piecewise linear approximation of a skin surface in Chapter 5. Finally, in Chapter 4 we extend the class of skin surfaces to the much wider class of envelope surfaces, in particular envelopes of balls.

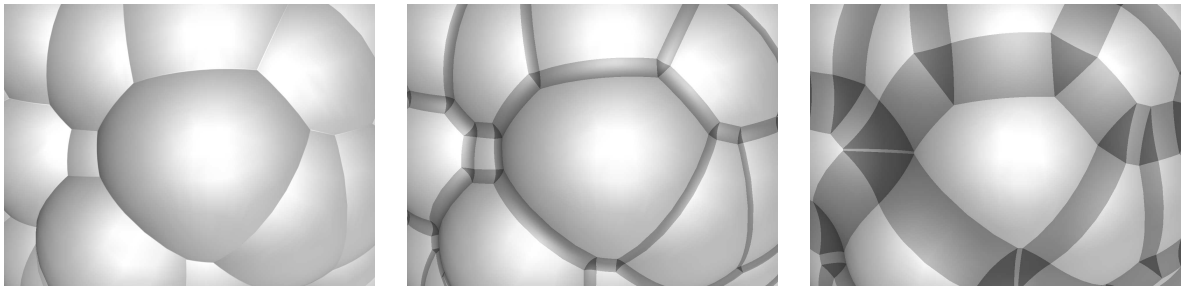

Figure 1.1: The skin surface blends spheres together

The remainder of the introduction is organized as follows. We start with an outline of this thesis. We then review areas in which surfaces emerge and present surface representations. In Section 1.4 we present surface models used for molecular surface design. We conclude the introduction with a brief description of several surface approximation algorithms.

1.1 Outline of the thesis

In this thesis we propose a class of surfaces represented by a finite set of balls contained inside the surface. The surface wraps tightly around the union of the balls, making the transition between balls smooth. The size of the smooth patches in between balls can be controlled as shown in Figure 1.1.

Surfaces defined by a set of balls are used in molecular modeling to build computer representations of large proteins. In [48] Edelsbrunner defines such a class of surfaces called *skin surfaces*. A skin surface is formed by a set of spheres that are 'blended' by smooth patches. These surfaces are used to simulate chemical processes, like protein folding [25], in which molecules deform. Several algorithms have been proposed to deform skin surfaces [28, 33, 48].

Generating balls Laser scanners provide an automated way to construct surfaces of physical objects. These scanners probe an object with a laser and produce a dense set of points on or near its surface. Many surface reconstruction algorithms have been proposed that compute a surface by reconstructing the object from these points, see Section 1.5. Algorithms [6, 7, 45] exist that transform a dense point sample on a surface into a set of balls almost contained inside the surface. We describe such an algorithm in Section 2.6. The generated balls can be used to define a surface in our representation.

An example of the generation of these balls is shown in Figure 1.2. The three pictures show the real-world object (the Stanford bunny), the point sample generated by a laser scanner and the constructed set of balls contained inside the bunny.

Surface approximation Even if a surface representation has many properties, like ease of visualization and that is suitable for further geometric processing, it is

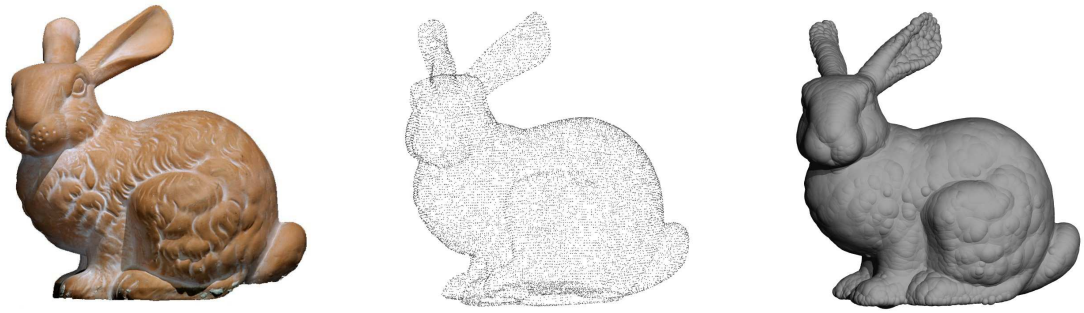

(a) The Stanford bunny

(b) A points on the surface of the bunny

(c) A set of balls contained in the bunny

Figure 1.2: The input for our algorithm

of little use if there is no (automatic) way to construct or approximate a surface in that representation. If the representation is not exact, then original surface and its approximation should at least be 'similar'. For example, one could say that two surfaces are similar if they are topologically equivalent, which, among other things, means that they have the same number of connected components, holes and tunnels. Another metric on the class of compact surfaces is the Hausdorff distance, which measures of the distance between the two surfaces. One could also require that not only the distance between the surfaces is small, but also that the normals to the surfaces agree. We discuss these notions in more detail in Section 2.2.

In this thesis, we use an additional criterion for good approximation. Our surfaces are represented by a set of balls. We assume that the boundary of their union forms a good approximation of the surface. Therefore, we also require that the approximating surface touches all input balls. This is a guarantee that the surface follows the boundary of the union of the balls closely.

1.1.1 Approximation by skin surfaces

We propose a method for the construction of a skin surface from an arbitrary surface that guarantees good approximation qualities.

A set of inscribed balls of a surface is dense if the boundary of the union of these balls and the surface have the same topology and a small Hausdorff distance. To become more precise, the boundary of the union of a dense set of balls is not differentiable on an intersection curve where two balls touch. To remove the tangent discontinuities, we use skin surfaces. A skin surface defines a smooth surface controlled by a single parameter called the shrink factor s. For a shrink factor equal to one, the skin surface is the boundary of the union of the balls. If the shrink factor is smaller than one, the radii of the balls are shrunk with a factor $\sqrt{s}$ and the balls are connected by smooth patches of hyperboloids and spheres. In order to make the

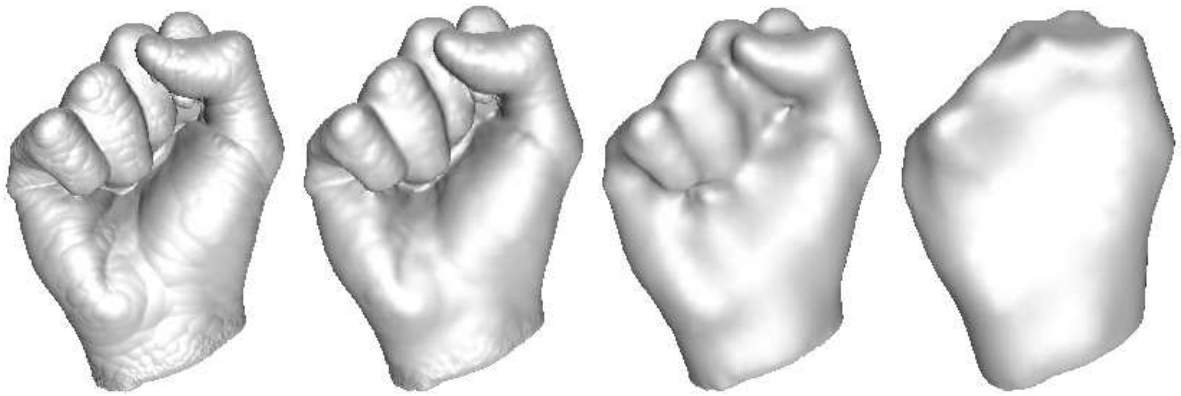

Figure 1.3: The approximating skin surface of a set of balls forming a hand. The patches between the balls are controlled by a single parameter.

skin surface wrap around the union of the balls (and not the shrunk balls), we first grow the input balls by a factor of $1/\sqrt{s}$. We call the skin surface with the grown balls the *extended skin surface*. An example of an extended skin surface is shown in Figure 1.3. The left figure shows the input balls and the shrink factor is decreasing from left to right.

If the shrink factor is close to one, the extended skin surface and the boundary of the union of the balls have small Hausdorff distance and the same topology. In fact, they are equal for a shrink factor one. As the shrink factor decreases, the patches become larger and they may cover (small) intrusions in the surface, viz. the region between the fingers and the palm in the third picture of Figure 1.3. Even different parts of the surface may grow towards each other and merge. Therefore it is important to have a lower bound on the shrink factor for which the extended skin surface and the union of the balls have the same topology and a Hausdorff distance smaller than some fixed constant. We propose an algorithm that derives the lower bound in Chapter 3.

By continuously decreasing the shrink factor, we are able to prove the following properties of the extended skin surface:

- the extended skin surface is topologically equivalent with the boundary of the union of the input balls;
- the Hausdorff distance between the union of the input balls and the body of the extended skin surface is smaller than some fixed constant $\epsilon > 0$;
- each ball in the input set is tangent to the extended skin surface.

In fact, for a dense set of balls, these results also apply to the surface from which the input balls are obtained if each ball contributes to boundary of their union.

1.1.2 Envelope surfaces

The extended skin surfaces have nice properties as mentioned in the previous section. The main disadvantage of the approach is that a skin surface is only controlled by

a single parameter, namely the shrink factor. For example in Figure 1.3, combining the fingers of the second picture and the palm and wrist of the third picture gives an approximation with the same properties as mentioned above.

In Chapter 4 we introduce a class of surfaces, called envelope surfaces, that forms a generalization of the class of skin surfaces. With envelope surfaces we first construct a triangulation of the centers of the balls and can change the envelope surface at the level of simplices of the triangulation yielding an approximation that is *locally adaptive*. Further, the interpolating patches generated by skin surfaces are always concave. For envelope surfaces these patches can be both concave and convex and, hence, allow for a better approximation. A two-dimensional example is given in Figure 1.4. The set of input circles is shown in Figure 1.4(a). Figure 1.4(b) shows an extended skin curve with a shrink factor close to one. The extended skin surface does not differ much from the boundary of the union of the input circles. For a smaller shrink factor patches between different fingers arise and the extended skin curve changes topology, as is shown in 1.4(c). Using an envelope surface to construct the surface, it is possible to interpolate in the direction of the fingers, but not in between the fingers, cf. Figure 1.4(d),

An envelope surface is defined as the boundary of the union of an infinite set of balls. We define these balls by constructing a continuous function that assigns a squared radius (weight) to each center. Under certain conditions on this weight function the envelope surface is tangent continuous. As an important special case, we analyze envelope surfaces defined by a piecewise quadratic weight function, which includes the class of skin surfaces. For this type of weight function, we develop efficient tests to validate the condition under which the envelope surface is tangent continuous.

Many desirable properties of skin surfaces hold in the generalized setting of envelope surfaces, which makes us believe that applications, like deformation schemes and meshing algorithms, developed for skin surfaces can be extended to envelope surfaces. For example, both the skin surface and the envelope surface are piecewise quadratic. However, a skin surface consists of parts of spheres and hyperboloids whereas the envelope surface consists of pieces of quadrics any type. Further, for both classes of surfaces there is a polyhedral partitioning of space decomposing the surface into these quadrics.

1.1.3 Meshing skin surfaces

Meshing of smooth surfaces is a first step in many numerical simulations and is often needed for visualization. In Chapter 5 we present an algorithm for meshing skin surfaces with guaranteed topology. The algorithm decomposes the skin surface into parts homeomorphic to a disk. A coarse mesh is then constructed by approximating each part with a piecewise linear patch.

The coarse mesh is not directly suitable for further processing since it has too little detail and many skinny triangles. Therefore we use several existing refinement algorithms to enhance the mesh. The first method we propose is a subdivision algorithm, which refines the mesh according to a subdivision scheme [69]. This

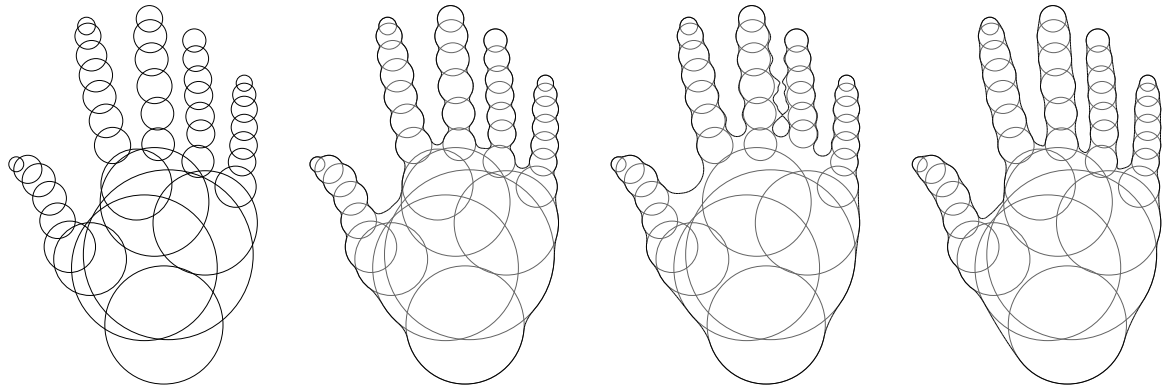

(a) An approximating set of circles

(b) The extended skin curve

(c) The extended skin curve with a lower shrink factor

(d) An envelope curve

Figure 1.4: Various interpolation schemes

algorithm is very fast and particularly useful for visualization of skin surfaces. A second algorithm is based on an algorithm by Chew [34] and increases the minimal angle of a triangle. This is necessary for the stability of numerical simulations, [23].

1.2 Generating surfaces

We discuss three areas in which surfaces emerge.

In the setting of *surface reconstruction* the surface is not completely known. A surface is constructed that forms a reconstruction of the original surface. For example, a laser scanner generates a finite sample of points lying on the surface. Based on that point sample, surface reconstruction algorithms construct a surface that is a reconstruction of the scanned object.

In Section 1.3, we see that the implementation of operations on a surface depend on its representation. Therefore it is sometimes necessary to change the representation of a surface. In the setting of *surface approximation*, a surface defined in one representation is approximated with a surface in another representation. Meshing algorithms for smooth surfaces form an important class of surface approximation method. A meshing algorithm generates a piecewise linear approximation (polyhedron) of a smooth surface. In general, the faces of the approximation are triangles. These meshes are used for visualizing the surface and for further geometric processing. The algorithms proposed in Chapter 3 and 5 belong to the class of surface approximation algorithms.

The last area from which surfaces emerge that we mention is *surface design*. Surface design is the process in which persons like engineers and architects construct a surface interactively using a computer. Important in surface design is that the surface can be easily and intuitively manipulated. The class of envelope surfaces,

introduced in Chapter 4 falls in this category.

1.3 Surface representations

We mention three basic classes of surface representations.

The first class is formed by *parametric surfaces* [52]. These surfaces are defined by a function that maps a two-dimensional domain to a surface in $\mathbb{R}^3$. An example of a parametric function $f : [0, \pi] \times [0, 2\pi] \to \mathbb{R}^3$ that defines the unit sphere uses spherical coordinates: $f(u, v) = (\sin(u)\sin(v), \sin(u)\cos(v), \cos(u))$.

It is easy to generate points on a surface defined by a parameterization, which is a starting point of many algorithms. These points are necessary for meshing surfaces. On the other hand it is hard to test whether a point lies on a parametric surface. Therefore it is also difficult to prove in general that the surface does not have self-intersections, i.e., that there are no two different parameter values that map to the same surface point.

The class of *implicit surfaces* [18] defines a surface as the zero set of a scalar valued function on $\mathbb{R}^3$. In implicit representation, the unit sphere is defined by the function $f : \mathbb{R}^3 \to \mathbb{R}$ given by $f(x) = \|x\|^2 - 1$. Testing whether a point lies on an implicit surface is easy. On the other hand, generating points on the surface is harder than for parametric surfaces. In fact, it is even hard to find the number of connected components and points on each connected component. Both the skin surface and the envelope surface are implicit surfaces.

Piecewise linear surfaces [49] play a crucial role in visualization and in numerical simulations. The representation of parametric surfaces and implicit surfaces are often too complicated to use directly in applications, therefore the surfaces are first meshed and the computations are then performed on the piecewise linear approximation. For further geometric processing it is important that the surface and the linear approximation are topologically equivalent. In Chapter 5, we construct a mesh from a skin surface with this property.

A surface representation that also defines a surface by a set of interior balls is the *Medial Axis Transform* (MAT) [56, 95, 99], see Section 2.6. The surface is defined as the boundary of the union of an infinite set of maximal balls. The centers of the balls form a skeletal structure of the surface called the medial axis. The balls are defined by assigning a radius to each point of the medial axis.

The thesis of Vermeer [95] focusses on the construction of a parametric representation of a surface from its Medial Axis Transform. They also analyze under which conditions on the medial axis and the radii of the balls the surface is tangent continuous. The methods we propose differ from the methods in [95] because we construct balls centered on a three-dimensional domain and not restricted to the medial axis. They also start with a complete representation of the MAT and then construct a parametric representation of the surface, whereas we assume that the we are given an incomplete set of balls, from which we reconstruct the surface. An advantage of our method over theirs is that the analysis under which conditions the boundary of the union of the balls is tangent continuous becomes easier. Our condition only

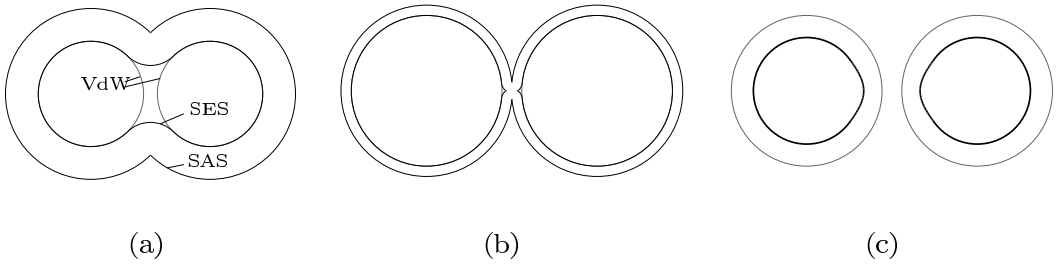

Figure 1.5: Different kinds of molecular surfaces. Figure (a) and (b) show the Van der Waals surface (VdW), the solvent accessible surface (SAS) and the solvent excluded surface (SES) for different sizes of the solvent probe. Figure (c) shows the skin surface.

depends on the radius function while their conditions also rely on the local structure of the medial axis.

1.4 Molecular modeling

The surfaces discussed in this thesis are related to skin surfaces. We present the theory of skin surfaces in more detail in Section 2.7. Several other types of surfaces have been proposed to model molecules. Some of these models give insight in the structure of the molecule or in characteristics like the double helix structure of DNA. To sketch the context in which skin surfaces were developed for molecular modeling, we review some of the models that are related to skin surfaces. A two-dimensional example of the different surfaces is shown in Figure 1.5.

Van der Waals surface Each atom in a molecule has an electrical charge that produces a Van der Waals force. This force is strongly repulsive in close proximity, becomes mildly attractive at intermediate range, and vanishes at long distance. Due to the repulsion on the short range, atoms cannot be too close to each other. The strength of the force depends on the type of atom. For each type of atom, the radius of the ball (the Van der Waals radius) where the center of no other atom can penetrate is experimentally determined. In this way atoms are associated with balls. The Van der Waals surface is the boundary of the union of these balls.

Solvent accessible surface Small gaps and pockets in the Van der Waals surface are not relevant for chemical processes [40, 77]. For example, if these gaps and pockets are so small that not even a solvent molecule (like water) can enter, they do not influence chemical processes. The solvent accessible surface is proposed as a tool to remove these small artifacts from the Van der Waals-surface.

To construct this surface, a spherical probe is rolled over the Van der Waals surface. The trace of the center of the probe is the solvent accessible surface. It bounds the region that is accessible by the (center of the) solvent. This surface is

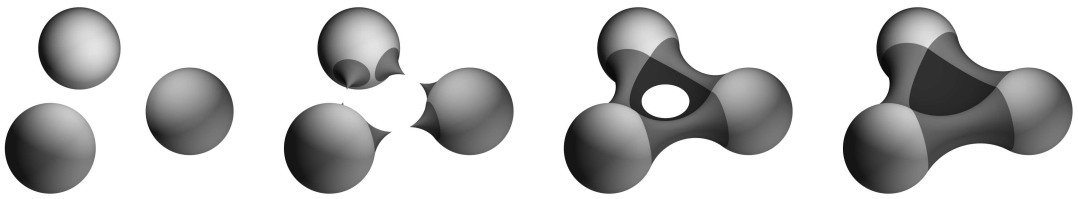

Figure 1.6: Solvent excluded surfaces for probe spheres with different radii. In light gray, the atoms, darker gray the interpolating patches between two balls and in dark gray the patches between three balls. The patches may not be tangent continuous as is the case in the second figure, where the tangent plane is discontinuous at six points. In the third figure the tangent plane of the SES is discontinuous at the intersection circle of the two probes that touch all three atoms.

equal to the boundary of the union of the balls obtained by adding the radius of the probe to the Van der Waals radius of each atom.

Solvent excluded surface Opposed to the solvent accessible surface, that bounds the region where the center of the probe enters, the solvent excluded surface [87] bounds the region where the solvent cannot enter. To construct this surface, the spherical probe is used to carve away the space outside the Van der Waals balls. This surface consists of spherical parts of the Van der Waals surface, toroidal patches between two Van der Waals balls and spherical patches where the probe touches three Van der Waals balls. This surface is also referred to as the Connelly surface [38, 39]. Note that this surface may also be not differentiable as shown in Figure 1.5(b) and Figure 1.6. Visualization algorithms for this type of surfaces are presented in [14, 15].

1.5 Surface reconstruction

Since the introduction of laser scanners, dense point sets of scanned objects are generally available. Surface reconstruction algorithms reconstruct the surface bounding the object from these point samples.

Many algorithms were developed that construct an *interpolating* or *approximating* surface from the set of points. An interpolating surface is a surface that passes through all sample points. Approximating surfaces do not interpolate the sample points, but lie in their vicinity. If the point sample contains noise, the points do not lie on the surface and algorithms that construct approximating surfaces can be used to smooth away the noise at the expense of a loss in detail.

We describe some algorithms for the purpose of surface reconstruction from point samples. This list is far from complete, but is intended to to give an overview of surface reconstruction algorithms related to this thesis. Most of these algorithms have theoretical guarantees with respect to the input surface like topological equality and a bound on the Hausdorff distance under the assumption that the point sample is an ϵ-sample, with $\epsilon < 0.1$. An ϵ-sample means that the point sample contains points on all parts of the surface and more points in regions with more detail.

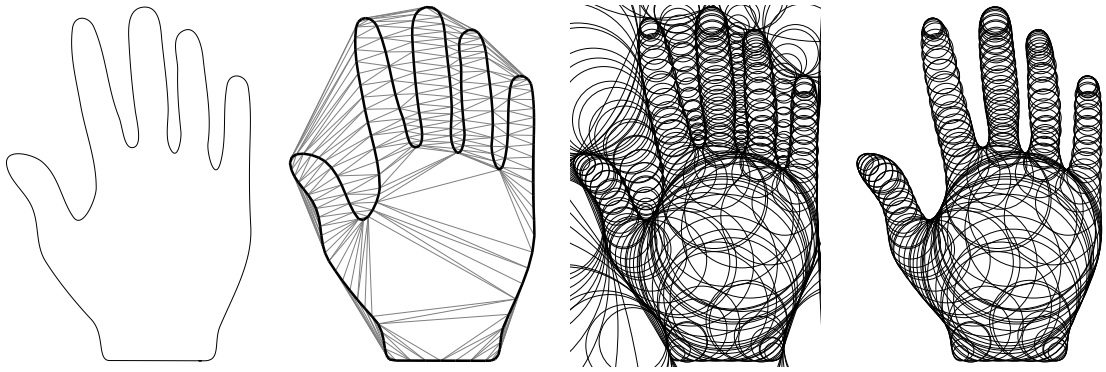

(a) The original curve

(b) The Delaunay triangulation of a finite point sample

(c) All polar circles

(d) The polar circles centered within the curve

Figure 1.7: Surface reconstruction based on Delaunay the triangulation.

A mathematical definition of the notion of an ϵ-sample is given in Section 2.6. Although it is hard to satisfy or test this condition, the algorithms work well in practice even if the sampling condition is not met.

Notation The set of sample points is denoted by $P = \{p_0, \ldots, p_n\}$. Some of the methods mentioned below produce an implicit function, the zero set of which is the reconstructed surface. These method of surface reconstruction are also called scattered data interpolation. We denote this function by F and the specified function value of F at the point p_i by f_i.

Delaunay based reconstruction Several algorithms reconstruct a piecewise linear surface using the Delaunay triangulation of the point set [5, 6]. These algorithms extend earlier work on curve reconstruction [4].

The Delaunay triangulation is one of the most basic structures in Computational Geometry, and is described in detail in Section 2.5. The Delaunay triangulation is the dual of the Voronoi diagram, which decomposes $\mathbb{R}^3$ into regions closest to a certain site. The farthest point in the Voronoi cell from a site is called the pole of the site. If the point sample is dense, the normal to the surface is well approximated by the vector from a site to its pole. A polar ball is the largest ball centered at a pole, not containing any site. Each polar ball has four sample points on its boundary.

The power crust algorithm [6] labels polar balls inside and outside. Two polar balls are neighbors if their Voronoi vertices are connected by a Voronoi edge. The algorithm starts by giving a large sphere outside the surface the label outside. Then it progresses to neighboring polar balls and gives these balls the same label if the outer angle in which the spheres intersect is large and the opposite label if the angle is small. The reconstructed surface is defined as the set of points with equal distance

to interior and exterior balls. Using a suitable distance function, this surface is a subset of the facets of the Delaunay triangulation, which is piecewise linear.

The different steps of the algorithm are depicted in Figure 1.7. The reconstructed curve is the subset of the edges of the Delaunay triangulation that are emphasized in Figure 1.7(b). We describe this algorithm in greater detail in Section 2.6.

The cocone algorithm [5] uses the poles to estimate the surface normals in the sample points. It then constructs cocones around the samples. These cocones are the complement of the cone with apex at the sample and symmetry axis in the direction of the estimated normal. The Delaunay facets (triangles) lying in the cocone are selected as candidate samples. A manifold extraction step selects triangles that form the final reconstructed surface.

Moving least squares The idea of the moving least squares approach [1, 79, 78] comes from the well known least squares technique to fit a surface to a set of points. The term "moving" refers to the various weighting of the points in calculating their contributions to the solution at different locations. In general, sites lying close have a larger weight. Several slightly different definitions are given in literature. for assigning the weights and constructing the moving least squares function.

The approximating MLS-surface is homeomorphic to the original surface if the uniform point sample is dense enough, as is shown in [70]. In [42], the condition on the uniform sample is made adaptive, hence requiring less points. In [90], Shen et.al. describe an algorithm that takes polyhedra as input instead of points.

Subdivision surfaces An important way of modeling and reconstructing surfaces is by means of subdivision surfaces. These surfaces are defined by a coarse control mesh which is then subdivided. Several schemes are proposed for subdivision. The schemes in [47, 100, 69] interpolate the vertices of the control mesh, whereas the schemes [81, 26] are approximating. In practice, the latter schemes produce better results.

In [65, 67, 66], Hoppe et al. describe a surface reconstruction algorithm that constructs this coarse mesh from an unorganized point cloud. The idea is to estimate the tangent plane of the surface in each sample points. This is done by a least squares fit through the samples in the neighborhood of the sample point. Then the normals are oriented consistently and the surface is extracted with a marching cubes [82] like algorithm.

This surface reconstruction algorithm is one of the first and very little is proved about their behavior.

Natural neighbors A natural neighbor surface [20, 53] constructs an implicit surface using natural neighbor coordinates, also called Sibson coordinates [91, 97]. The implicit surface is the zero-set of the implicit function.

The definition of natural neighbor coordinates is given in terms of Voronoi cells, which are more generally discussed in Section 2.5. In short, the Voronoi cell ν_p of a

point $p \in P$ with respect to a point sample P is the set of points closer to p than to any other sample in P.

The natural neighbor coordinate associated to P of a point $x \in \text{conv}(P)$ are the functions

$$\sigma_i(x) = \frac{|\nu_i(x)|}{|\nu_x|},$$

where $|\cdot|$ denotes the volume of a set, ν_x the Voronoi cell of the point x with respect to $P \cup \{x\}$ and $\nu_i(x)$ the intersection of ν_x with the Voronoi cell of $p_i \in P$ with respect to P. Since the Voronoi diagram partitions $\mathbb{R}^3$, it follows that $\sum \sigma_i(x) = 1$, for all $x \in \mathbb{R}^3$. Further, $\sigma_i(p_j) = \delta_{ij}$, with δ_{ij} the Kronecker delta. The points p_i with $\sigma_i(x) > 0$ are called natural neighbors of x.

The interpolated function value $F(x)$ is defined as

$$F(x) = \sum_i \sigma_i(x) f_i,$$

or, if the normal n_i of site p_i is known, as

$$F(x) = \sum_i \sigma_i(x)(f_i + \frac{1}{2}\langle n_i, x - p_i\rangle).$$

The approximation is the zero-set of F. The function F is differentiable, except at sample points.

Initially, the natural neighbor coordinates were proposed for function interpolation. For surface reconstruction, only natural neighbors lying near an estimated tangent plane in the query point x are used. These points are called the T-neighbors.

Radial basis functions The reconstruction algorithms that use radial basis functions [86, 93, 94] construct an implicit function F, the zero set of which is the reconstructed surface. The implicit function is a weighted sum of radially symmetric functions $\phi : \mathbb{R}^+ \to \mathbb{R}$, depending only on the distance to a site:

$$F(x) = \sum_{i=0}^{n} w_i \, \phi(\|p_i - x\|) + P(x), \tag{1.1}$$

where w_i is the weight of the radial basis function ϕ centered at p_i and P is (an optional) low degree polynomial. Common choices are $\phi(r) = r$ (linear), $\phi(r) = r^2 \log(r)$ (thin-plate spline or biharmonic radial basis function in 2D) and $\phi(r) = r^3$ (thin-plate spline in 3D).

In Equation (1.1), the unknowns are the weights w_i and the coefficients c_i of the polynomial P. Since the equation is linear in both parameters, it can be expressed as a linear system. The weights w_i are used to solve the equations $F(p_i) = f_i$ and the coefficients of P are used to guarantee that the system of equations has a solution.

These equations are formulated in the following linear system:

$$\begin{pmatrix} \phi_{01} & \phi_{02} & \dots & \phi_{0n} & p_{0,0} & \dots & p_{0,m} \\ \phi_{11} & \phi_{12} & \dots & \phi_{1n} & p_{1,0} & \dots & p_{1,m} \\ \vdots & \vdots & & \vdots & \vdots & & \vdots \\ \phi_{n1} & \phi_{n2} & \dots & \phi_{nn} & p_{n,0} & \dots & p_{n,m} \\ & & & \text{constraints induced by polynomial} & & & \end{pmatrix} \begin{pmatrix} w_0 \\ w_1 \\ \vdots \\ w_n \\ c_0 \\ \vdots \\ c_n \end{pmatrix} = \begin{pmatrix} f_0 \\ f_1 \\ \vdots \\ f_n \\ 0 \\ \vdots \\ 0 \end{pmatrix},$$

where $\phi_{ij} = \phi(\| p_i - p_j \|)$ and $p_{i,k}$ is the factor of P with coefficient c_k evaluated at p_i.

If all function values f_i are zero, then the trivial solution $w_i = 0, c_j = 0$ is valid. In that case F is the zero function. To avoid this solution, additional samples are placed inside (outside) the surface with positive (negative) function values. The position of these additional constraints influences the surface. Common placement rules are along the normal (inside, outside or both) or on a skeletal structure of the surface called the medial axis which is described in Section 2.6.

Chapter 2

Preliminaries

This chapter introduces the necessary mathematical notions and definitions used throughout the thesis.

2.1 Topology

Topology is the mathematical study of the invariants that are preserved through deformations, without changing the local connectivity of the object. Topological invariants include the number of connected components, the number of holes and the number of tunnels. In the topological sense a doughnut and a teacup are equivalent, since both are made from one piece and both have one tunnel (the hole in the doughnut and the ear of the teacup).

Two objects $X_0, X_1 \subset \mathbb{R}^d$ are *homeomorphic* or *topologically equivalent* if there is a continuous bijective map $f : X_0 \to X_1$ that has a continuous inverse. The two circles in Figure 2.1(a) and (b) are homeomorphic.

A (topological) embedding is a continuous injective map, which is a homeomorphism onto its image. Let $X_0, X_1 \subset \mathbb{R}^d$ and $f_0, f_1 : Z \to \mathbb{R}^d$ be two embeddings such that $f_i(Z) = X_i$. Then X_0 and X_1 are *isotopic* if there is a continuous family F_s, with $0 \leq s \leq 1$, of embeddings $Z \to \mathbb{R}^d$, such that $F_0 = f_0$ and $F_1 = f_1$.

In particular, isotopic surfaces are homeomorphic. The two circles in Figure 2.1(a) and (b) are not isotopic.

2.2 Metrics for surface approximation

In surface approximation, one surface is approximated as good as possible with another surface. However it is not clear what 'good' exactly means. A metric on a topological space X is a non-negative function g with the following properties:

- $g(x, y) \geq 0$ and $g(x, y) = 0$ if $x = y$, with $x, y \in X$;
- $g(x, y) = g(y, x)$ (a metric is symmetric);
- $g(x, y) + g(y, z) \geq g(x, z)$ (the triangle inequality).

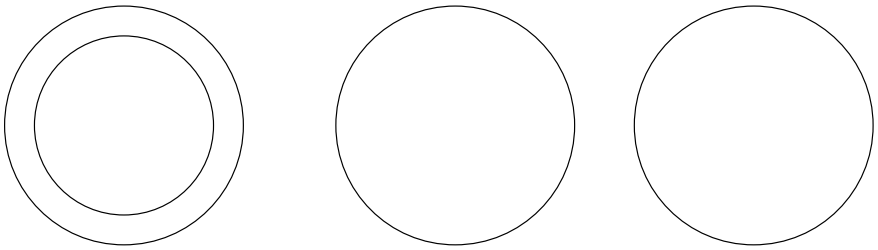

(a) (b)

Figure 2.1: The circles in (a) and (b) are homeomorphic, but not isotopic.

Capturing the right topology on its own is not enough to guarantee a good approximation. The Hausdorff distance is a measure for the distance between two sets (the surface is an infinite set of points), see [60, Chapter 43]. The Hausdorff distance $d(X, Y)$ of two compact subsets X, Y is defined by:

$$d(X, Y) = \max(\max_{x \in X} \min_{y \in Y} g(x, y), \max_{y \in Y} \min_{x \in X} g(x, y)).$$

Note that $\max_{x \in X} \min_{y \in Y} \| x - y \|$ and $\max_{y \in Y} \min_{x \in X} \| x - y \|$ are not necessarily equal. For example if X is contained in Y then $\max_{x \in X} \min_{y \in Y} \| x - y \|$ is zero, but $\max_{y \in Y} \min_{x \in X} \| x - y \|$ is positive.

2.3 Envelopes

Let $F : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}$ be a C^1 function defining a one parameter family of functions $F_\mu : \mathbb{R}^d \to \mathbb{R}$ by $F_\mu(x) = F(x, \mu)$. Assume that the sets

$$Z_\mu = \{x \in \mathbb{R}^d \mid F_\mu(x) \leq 0\}$$

are bounded. The *envelope* of the family Z_μ is the boundary of their union. Note that the envelope may be empty. Let $x \in Z_{\mu_0}$ be a point on the envelope. Then $F(x_0, \mu_0) = 0$, and $F(x_0, \mu) \geq 0$ for all μ. Therefore μ_0 is a (global) minimum of the function $\mu \mapsto F(x_0, \mu)$, so $\frac{\partial F}{\partial \mu}(x_0, \mu_0) = 0$. It follows that the envelope is a subset of the *discriminant set* D_F of the family F, defined by

$$D_F = \{x \in \mathbb{R}^d \mid F(x, \mu) = \frac{\partial F}{\partial \mu}(x, \mu) = 0, \text{ for some } \mu \in \mathbb{R}\}.$$

If the family depends on a multi-dimensional parameter $\mu \in \mathbb{R}^k$ then the discriminant set is defined by

$$D_F = \{x \in \mathbb{R}^d \mid F(x, \mu) = 0, \ \nabla_\mu F(x, \mu) = 0, \text{ for some } \mu \in \mathbb{R}^k\}.$$

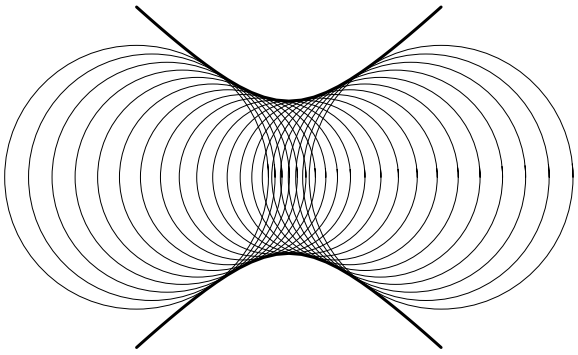

Figure 2.2: The envelope of the circles given by $F((\xi_1, \xi_2), \mu) = (\xi_1 - \mu)^2 + \xi_2^2 - 0.5\mu^2 - 0.5$.

Example We compute the envelope of a set of spheres centered at $(\mu, 0, 0)$ and with squared radius $s\mu^2 + c$. The function $F : \mathbb{R}^3 \times \mathbb{R} \to \mathbb{R}$ defines this set of spheres:

$$F(x, \mu) = (\xi_1 - \mu)^2 + \xi_2^2 + \xi_3^2 - s\mu^2 - c,$$

with $x = (\xi_1, \xi_2, \xi_3)$. Hence $\frac{\partial F}{\partial \mu}(x, \mu) = 2((1 - s)\mu - \xi_1)$, which vanishes if $\mu = \xi_1/(1 - s)$. Substitution of μ in F yields an implicit equation for the envelope:

$$-\frac{s}{1 - s}\xi_1^2 + \xi_2^2 + \xi_3^2 - c = 0.$$

In Chapter 4 we analyze envelopes of spheres with a piecewise quadratic radius-squared function in more detail.

2.4 Complexes

A *convex polyhedron* is the, possibly unbounded, intersection of a finite number of closed half-spaces. A *polyhedral complex* $\mathcal{C}$ is the collection of faces of a finite number of polyhedra, such that

(i) if $\sigma \in \mathcal{C}$ and τ is a face of σ, then $\tau \in \mathcal{C}$ and

(ii) if $\sigma, \sigma' \in \mathcal{C}$ such that $\sigma \cap \sigma' \neq \emptyset$, then $\sigma \cap \sigma'$ is a face of both.

We write $\tau \leq \sigma$, for $\tau, \sigma \in \mathcal{C}$, if τ is a face of σ. A *d-simplex* is the convex hull of $d + 1$ independent points. A *simplicial complex* is a polyhedral complex all non-empty faces of which are simplices.

2.5 Voronoi diagrams and Delaunay triangulations

The Voronoi diagram, named after Georges Voronoi [46, 96], is one of the most fundamental structures in Computational Geometry. You may think of the Voronoi diagram in terms of a *grass-fire analogy*. At each site a fire is ignited that spreads

with constant rate of one unit distance per unit time. The Voronoi diagram consists of cells that are ignited by the same site (or set of sites).

The Voronoi diagram has applications in many areas of interest. Different names particular to the respective field have been used such as medial axis transform in biology and physiology, Wigner-Seitz zones in chemistry and physics, domains of action in crystallography, and Thiessen polygons in metereology and geography,[13]. For a survey on Voronoi diagrams and its applications see [12, 54] and [60, Chapter 23].

More formally, the Voronoi diagram of a collection of geometric objects $\mathcal{P}$ is a partition of space into cells, each of which consists of the points closer to one particular object than to any other. For each site $p \in \mathcal{P}$, the *Voronoi cell* ν_p is defined by:

$$\nu_p = \{x \in \mathbb{R}^d \mid d(x,p) \leq d(x,p'), \text{ for all } p' \in \mathcal{P}\}.$$

The input sites are not necessarily points. In $\mathbb{R}^2$, Voronoi diagrams of geometrical objects like line segments, polygons and circles have been studied [2, 64, 83].

Different distance functions yield different Voronoi diagrams. Usually, the Euclidean distance or two-norm (L_2) denoted by $\|\cdot\|$ is chosen as the distance function d.

By giving sites a weight it is possible to favor one site over another. A weighted point $\hat{p} = (p, P) \in \mathbb{R}^d \times \mathbb{R}$ has a location p and a weight P. A point with zero weight is called an *unweighted point*. The way this weight is taken into account results in different Voronoi diagrams. In the additively weighted Voronoi diagram the distance function d is given by $d(\hat{p}, x) = \|p - x\| - P$. The multiplicatively weighted Voronoi diagram is defined by the distance function given by $d(\hat{p}, x) = P\,\|p - x\|$.

Power diagrams The last, and the distance function used in this thesis is the power distance $\pi(\hat{p}, x) = \|p - x\|^2 - P$, where $\|\cdot\|$ denotes the Euclidean distance. The Voronoi diagram obtained from this distance function is the power diagram. If all weights are equal, then the power diagram is just the unweighted, Euclidean Voronoi diagram. Since the power diagram is the only type of Voronoi diagram used in this thesis, we also refer to it as the (weighted) Voronoi diagram.

The power distance is generalized to a pseudo distance between two weighted points $\hat{p}$ and $\hat{q}$:

$$\pi(\hat{p}, \hat{q}) = \|p - q\|^2 - P - Q, \tag{2.1}$$

where $\hat{p} = (p, P)$, $\hat{q} = (q, Q)$. The unweighted points with zero power distance to a weighted point $\hat{p}$ form the sphere centered at p with radius $\sqrt{P}$. We make no distinction between the weighted point and this sphere. If the weight is negative, the sphere is imaginary.

Note that the power distance π is not a real distance function since $\pi(\hat{p}, \hat{q})$ is not positive nor does it satisfy the triangle inequality or is $\pi(\hat{p}, \hat{q})$ zero if $\hat{p} = \hat{q}$.

Two weighted points with zero power distance are called *orthogonal*. If the weights of the two points are positive, this means that the corresponding spheres intersect in a circle and their tangent planes are perpendicular. Two weighted points are *further than orthogonal* if their power distance is positive. An *orthosphere* of a

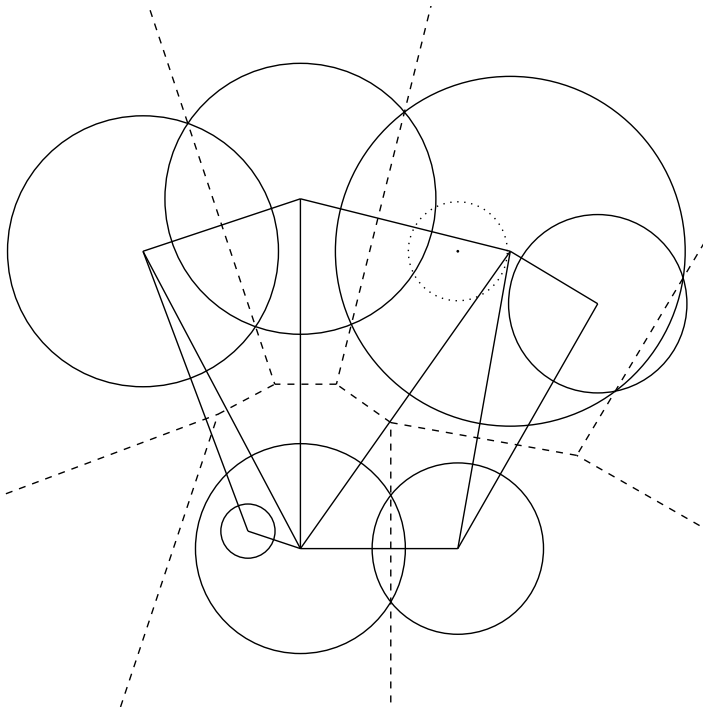

Figure 2.3: The Voronoi diagram (dashed) and the Delaunay triangulation (solid) of a set of weighted points (circles) in 2D. The dotted circle is hidden.

set of weighted points $\mathcal{P}$ is, by definition, a sphere orthogonal to each of the weighted points in $\mathcal{P}$. An unique orthosphere exists for $d+1$ weighted points in $\mathbb{R}^d$ that are in general position.

Let $\hat{p}$ and $\hat{q}$ be two weighted points. The points with equal power distance to $\hat{p}$ and $\hat{q}$ form a hyperplane orthogonal to the line segment pq. Denote with $H_{\hat{p},\hat{q}}$ the halfspace containing all points closer to $\hat{p}$ than $\hat{q}$. By definition, the weighted Voronoi cell $\nu_{\hat{p}}$ of the site $\hat{p}$ is the intersection of a set of halfspaces which is a, possibly unbounded, convex polyhedral region:

$$\nu_{\hat{p}} = \bigcap_{\hat{q}\in\mathcal{P},\hat{q}\neq\hat{p}} H_{\hat{p},\hat{q}}.$$

The weighted Voronoi cell is unbounded if the center of the weighted point lies on the convex hull of the centers of all weighted points in the input set. By choosing appropriate weights, a weighted Voronoi cell can also be empty. A weighted point with an empty Voronoi cell is called *hidden*, or redundant. A Voronoi cell of the power diagram also does not necessarily contain the location of its site.

The two-dimensional power diagram of eight weighted points, of which one is hidden (the dotted circle), is shown in Figure 2.3. The power diagram is drawn with dashed lines and the solid lines depict its dual, viz. the Delaunay triangulation.

The Voronoi diagram inherits a polyhedral structure by generalizing the definition of a Voronoi cell to sets of sites $\mathcal{X} \subseteq \mathcal{P}$:

$$\nu_{\mathcal{X}} = \{x \in \mathbb{R}^d \mid \pi(x,p) \leq \pi(x,p'), \text{ for all } p \in \mathcal{X}, p' \in \mathcal{P}\}$$

In other words, the Voronoi cell $\nu_{\mathcal{X}}$ is the intersection of the Voronoi cells $\nu_{\hat{p}}$ for $\hat{p} \in \mathcal{X}$. Hence, $\nu_{\mathcal{X}}$ is contained in $\nu_{\mathcal{X}'}$ if $\mathcal{X}' \subseteq \mathcal{X}$.

General position In this thesis we assume general position, by which we mean that no 5 weighted points are equidistant to a point in $\mathbb{R}^3$ and no $k+2$ centers of weighted points lie on a common k-flat for $k = 0, 1, 2$. We also assume that no four spheres bounding balls in the input set touch in a single point. Several methods like [50] exist to symbolically perturb a data set and ensure these conditions. Also algorithms exist that actually perturb the input in such a way that the perturbed input is in general position [55, 61].

Under this genericity condition $\nu_{\mathcal{X}}$ is determined by $|\mathcal{X}|-1$ independent equalities (and a number of inequalities). Therefore $\dim \nu_{\mathcal{X}} = 3 - (|\mathcal{X}| - 1) = 4 - |\mathcal{X}|$ if $\nu_{\mathcal{X}}$ is not empty. Hence, $|\mathcal{X}| \leq 4$ for a non-empty Voronoi cell. Every 0-cell is a point, which is the intersection of four adjacent Voronoi cells. A 1-cell is a line segment, possibly unbounded. And a facet (2-cell) is the intersection of the boundary of two 3-cells which is a convex polyhedron. The 2-cell is a subset of the set of unweighted points equidistant to the two weighted points associated with the 3-cells. A 3-cell is a nonempty Voronoi-cell of some $\hat{p} \in \mathcal{P}$. The Voronoi cell $\nu_{\mathcal{X}}$ for $\mathcal{X} \subseteq \mathcal{P}$ is non-empty, iff there is an orthosphere of $\mathcal{X}$ that is further than orthogonal from all weighted points in $\mathcal{P} \setminus \mathcal{X}$, cf. [49, Chapter 5].

To conclude, the *weighted Voronoi diagram* (power diagram) is the subdivision of 3-space generated by the non-empty weighted Voronoi cells:

$$\mathrm{Vor}(\mathcal{P}) = \{\nu_{\mathcal{X}} | \mathcal{X} \subseteq \mathcal{P} \wedge \nu_{\mathcal{X}} \neq \emptyset\}.$$

Decomposition of the union of balls The weighted Voronoi diagram decomposes the boundary of the union of a set of balls into spherical patches. A point on the intersection of a ball $\hat{p}$ and the boundary of the union of balls has zero distance to $\hat{p}$ and non-negative distance to all other balls. Therefore this point is contained in the Voronoi cell $\nu_{\hat{p}}$ and

$$(\partial \bigcup \mathcal{P}) \cap \hat{p} = \nu_{\hat{p}} \cap \partial \hat{p}.$$

We use this property to decompose the boundary of the union of a set of balls into spherical patches for the meshing algorithm proposed in Chapter 5.

2.5.1 Delaunay triangulations

The Delaunay triangulation is the dual of the Voronoi diagram, which means that there is a Delaunay simplex $\delta_{\mathcal{X}}$ for every non-empty Voronoi cell $\nu_{\mathcal{X}}$. The weighted Delaunay triangulation, or regular triangulation, is the dual of the weighted Voronoi diagram. The weighted Delaunay triangulation and the Delaunay triangulation are equal if all weights are equal since in that case the power diagram and the unweighted Voronoi diagram are identical.

The convex and affine hull of a set of (weighted) points $\mathcal{X} = \{\hat{p}_i\}$ is defined by:

$$\begin{aligned} \mathrm{aff}(\mathcal{X}) &= \{\sum_i \gamma_i \hat{p}_i \mid \sum_i \gamma_i = 1\} \\ \mathrm{conv}(\mathcal{X}) &= \{\sum_i \gamma_i \hat{p}_i \mid \sum_i \gamma_i = 1, \gamma_i \geq 0\} \end{aligned}$$

A *weighted Delaunay* $(3-\ell)$*-cell* exists for every non-empty weighted Voronoi ℓ-cell $\nu_{\mathcal{X}} \in \mathrm{Vor}(\mathcal{P})$ and is the convex hull of the centers of the weighted points in $\mathcal{X}$:

$$\delta_{\mathcal{X}} = \mathrm{conv}(\{p | \hat{p} \in \mathcal{X}\}).$$

Observation 1. *The cells $\delta_{\mathcal{X}}$ and $\nu_{\mathcal{X}}$ are complementary and orthogonal.*

The focus Note that $\delta_{\mathcal{X}}$ and $\nu_{\mathcal{X}}$ for $\nu_{\mathcal{X}} \neq \emptyset$ can be disjoint, but their affine hulls always intersect in a single point, the *focus* $f(\mathcal{X})$ of $\mathcal{X}$. Since the focus lies on the affine hull of the weighted Voronoi cell, it follows that $\pi(\hat{p}, f(\mathcal{X})) = \pi(\hat{q}, f(\mathcal{X}))$ for all $\hat{p}, \hat{q} \in \mathrm{conv}(\mathcal{X})$.

The *weighted Delaunay triangulation* is the subdivision generated by the Delaunay cells:

$$\mathrm{Del}(\mathcal{P}) = \{\delta_{\mathcal{X}} \mid \nu_{\mathcal{X}} \in \mathrm{Vor}(\mathcal{P})\}.$$

Geometric construction The definition of the weighted Delaunay triangulation given above is indirect in the sense that it is derived from the weighted Voronoi diagram. There is also a direct geometric construction of the weighted Delaunay triangulation. For this construction we introduce the vector space of weighted points, cf. [48]. The set of weighted points inherits a vector space structure from $\mathbb{R}^4$ via the bijective map $\Pi : \mathbb{R}^3 \times \mathbb{R} \to \mathbb{R}^4$, defined by

$$\Pi(\hat{p}) = (\xi_1, \xi_2, \xi_3, \|p\|^2 - p), \tag{2.2}$$

with $p = (\xi_1, \xi_2, \xi_3)$. Using this vector space we define the addition of two weighted points and the multiplication of a weighted point by a scalar as

$$\begin{aligned} \hat{p} + \hat{q} &= \Pi^{-1}(\Pi(\hat{p}) + \Pi(\hat{q})), \\ c \cdot \hat{p} &= \Pi^{-1}(c \cdot \Pi(\hat{p})). \end{aligned}$$

For a skin surface with a shrink factor s we introduce the notion of *shrinking*. Let $\hat{p}$ be a weighted point and $\hat{p}'$ the weighted point centered at p with zero weight. The shrunk weighted point $\hat{p}^s$ is defined as:

$$\hat{p}^s = s\,\hat{p} + (1-s)\hat{p}'. \tag{2.3}$$

If $\mathcal{P}$ is a set of weighted points, we denote by $\mathcal{P}^s$ the set obtained by shrinking every point of $\mathcal{P}$ by a factor s.

Before we state the geometric construction of the Delaunay triangulation, we make some preliminary observations on the vector space structure inherited under Π.

Let $\hat{p}' = (p', P')$ be an affine combination of the weighted points $\hat{p}_i$, i.e., $\hat{p} = \sum \gamma_i \hat{p}_i$ for scalars γ_i with $\sum \gamma_i = 1$

$$\hat{p}' = \Pi^{-1}\left(\sum_i \gamma_i \Pi(\hat{p}_i)\right),$$

A simple calculation shows that

$$\begin{aligned} p' &= \textstyle\sum_i \gamma_i p_i, \\ P' &= \|p'\|^2 - \textstyle\sum_i \gamma_i (\|p_i\|^2 - P_i) \end{aligned} \tag{2.4}$$

It follows from Equation (2.4) that $\hat{p}^s = (p, s\,P)$. In other words, $\hat{p}^s$ corresponds to the ball centered at p with radius $\sqrt{sP}$. We express the power distance of $\hat{p}'$ to an unweighted point x in terms of the power distance to the weighted points $\hat{p}_i$:

$$\begin{aligned} \pi(\hat{p}', x) &= \|x - p'\|^2 - P' \\ &= \|x - p'\|^2 - \|p'\|^2 + \sum_i \gamma_i(\|p_i\|^2 - P_i) \\ &= \|x\|^2 - 2\langle x, p'\rangle + \sum_i \gamma_i(\|p_i\|^2 - P_i) \\ &= \sum_i \gamma_i(\|x\|^2 - 2\langle x, p_i\rangle + \|p_i\|^2 - P_i) \\ &= \sum_i \gamma_i\, \pi(\hat{p}_i, x) \end{aligned} \tag{2.5}$$

Without changing the derivation, we may replace x with a weighted point $\hat{q}$.

Observation 2. *Let $\hat{p}'$ be an affine combination of weighted point $\hat{p}_i$. Then $\hat{p}'$ is orthogonal to the orthosphere of the weighted points $\hat{p}_i$.*

This is not true for affine combinations in the Euclidean space, but follows from the change of weights induced by the vector space inherited under Π.

Finally, note that if Π maps a weighted point $\hat{p}$ vertically above another weighted point $\hat{q}$, then their centers are identical and the weight of $\hat{p}$ is smaller than the weight of $\hat{q}$. Hence the weighted points on the lower convex hull of $\Pi(\mathcal{P})$ have maximal weight.

Lemma 3. *A set $\mathcal{X} \subseteq \mathcal{P}$ defines a Delaunay simplex $\delta_{\mathcal{X}}$ if* $\mathrm{conv}(\Pi(\mathcal{X}))$ *lies on the lower convex hull of* $\Pi(\mathcal{P})$.

Proof. Assume that the lifted weighted point in $\mathrm{conv}(\mathcal{X})$ lies on the lower convex hull of $\mathrm{conv}(\mathcal{P})$, then there exists a hyperplane $h_{\mathcal{X}}$ containing $\Pi(\mathrm{conv}(\mathcal{X}))$ and lying vertically below the convex hull of $\Pi(\mathcal{P})$. Let $\hat{q}$ be the orthosphere of all weighted points projecting onto $h_{\mathcal{X}}$. This orthosphere exists since it is defined by any four independent points on $h_{\mathcal{X}}$ and orthogonal to all points by Equation (2.5). Since $h_{\mathcal{X}}$ lies below the convex hull of $\Pi(\mathcal{P})$, it follows from the definition of Π that the weight

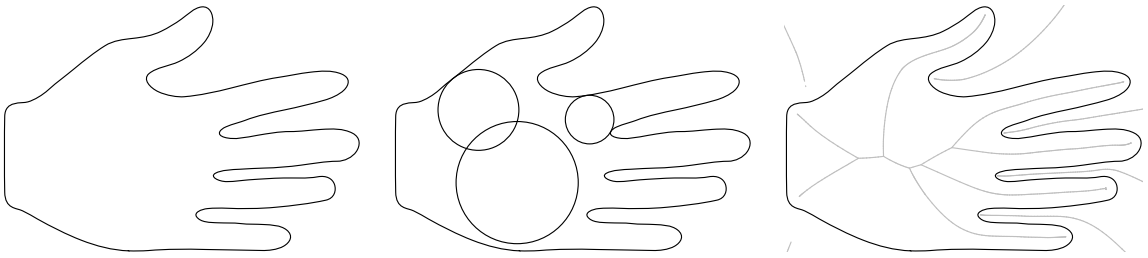

(a) A closed curve

(b) Some empty circles. The smallest circle is also maximal

(c) The medial axis inside and outside the curve

Figure 2.4: The medial axis in 2D.

of $\hat{\mathfrak{p}} \in \mathcal{P}$ is smaller than the weight of the weighted point $\hat{\mathfrak{p}}'$ lying on $\mathfrak{h}_{\mathcal{X}}$ with center $\mathfrak{p}$. Since $\hat{\mathfrak{p}}'$ and $\hat{\mathfrak{q}}$ are orthogonal, $\hat{\mathfrak{p}}$ is further than orthogonal from $\hat{\mathfrak{q}}$ and $\mathcal{X}$ forms a Delaunay simplex.

Conversely, if the lifted weighted point in $\text{conv}(\mathcal{X})$ does not lie on the lower convex hull of $\text{conv}(\mathcal{P})$, then no hyperplane $\mathfrak{h}_{\mathcal{X}}$ containing $\text{conv}(\mathcal{X})$ and lying below $\text{conv}(\mathcal{P})$ exists. Hence there is no orthosphere of $\mathcal{X}$ that is further than orthogonal to all weighted points in $\mathcal{P}$. □

2.6 Medial axis and the Medial axis transform

Much research has been done on constructing skeletal representations of two-dimensional curves and three-dimensional surfaces. Perhaps the most natural concept is the *medial axis* introduced in [19]. The medial axis of a surface is the closure of the set of points that have more than one closest point on the surface. A recent survey paper on the medial axis is [11]. The medial axis of a generic surface consists of surface patches that meet in a well defined way. For curves this structure is analyzed with differential geometry in [57, 59] and for surfaces in [58].

By definition, each point of the medial axis is the center of a ball touching the surface in at least two points and containing no part of the surface. An *empty ball* is a ball that is contained in the surface. *Maximal balls* are empty balls that are maximal with respect to inclusion. In other words, an empty ball is maximal if there is no empty ball in which it is contained. Each maximal ball touches the surface in at least two points and is therefore centered on the medial axis. The set of maximal balls of a surface is the *medial axis transform* (MAT). The boundary of the union of the maximal balls is the surface [98], which means that the medial axis transform is a representation of the surface. In fact, the medial axis is a strong deformation retract of the region bounded by the surface, see [35, 80]. The construction of the surface from a medial axis transform is studied in [56, 95].

A two-dimensional example depicting empty circles, maximal circles and the

medial axis is given in Figure 2.4. Figure 2.4(a) shows a closed curve and (b) the same curve with three empty balls. The only maximal circle in the figure is the smallest circle. The medial axis of the curve is shown in (c). Note that the outer branches are clipped, but are actually unbounded.

The medial axis and the medial axis transform are of interest in this thesis for two reasons. First, the union of the maximal balls contained in the interior of a closed surface define the surface from the inside. Although it is in general hard to compute the medial axis transform of a smooth surface, it is possible to construct a dense finite sample of approximate maximal balls from a dense point sample on the surface. An approximate maximal ball is a ball that has a small Hausdorff distance to a true maximal ball. We describe the construction of a sample of approximate maximal balls later in this section. The boundary of the union of the approximate maximal balls has a small Hausdorff distance to the original surface. This sample is therefore well suited as input for the algorithms described in Chapter 3 and 4.

Secondly, the distance to the medial axis is a good measure of the level of detail of the surface. For example, the distance to the medial axis in Figure 2.4 is smaller near the fingers than on the part of the curve forming the hand palm. This distance is called the *local feature size* (LFS). Using the local feature size we make mathematically precise what we mean with a dense point sample. Using existing algorithms, it is possible to construct the set of approximate maximal balls from a dense point sample.

Local feature size

The local feature size in a surface point is the distance from that point to the medial axis of the surface. It is a measure for the level of detail in a certain part of a surface. This measure is used to describe what is meant by a dense sample of points on a surface.

A point sample on the surface is an ϵ-sample if the distance in any point x on the surface to a sample is at most ϵ times the local feature size in x. For $\epsilon < 0.1$, the power crust algorithm [6] produces a set of balls such that the boundary of its union S' approximates the surface in the sense that:

- The Hausdorff distance between S and S' is of the order of ϵ, [6, Theorem 21].
- The angle between the normal in a regular point on S' and the normal in the closest point on S is of the order $O(\sqrt{\epsilon})$ [6, Theorem 24].
- The surfaces are homeomorphic [6, Lemma 25].

The balls contained in the bunny in Figure 1.2(c) are generated by the power crust algorithm.

Approximate maximal balls

An approximate maximal ball is a ball that has small Hausdorff distance to a true maximal ball. A set of approximate maximal balls can be constructed with the

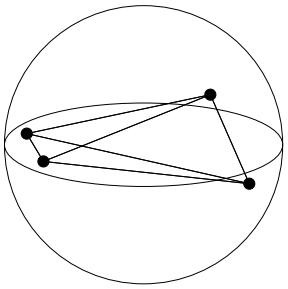

Figure 2.5: A sliver.

power crust algorithm [6, 7]. A similar algorithm is presented in [45]. The algorithm depends heavily on the Voronoi diagram of the point set. A point on a Voronoi face has at least two nearest sites. Hence, the two-skeleton of the Voronoi diagram is the medial axis of the point sample.

In 2D, the vertices of the Voronoi diagram are the points that have three closest sites. For an ϵ-sample of the curve, the vertices of the Voronoi diagram lie close to the medial axis. An example is shown in Figure 2.4. The crucial step is to label the Voronoi vertices as inside or outside. To this end, the algorithm constructs a weighted Voronoi diagram of the empty circles centered at Voronoi vertices. It gives an initial circle outside the surface the label outside. Then it progresses to neighboring empty circles and gives these circles the same label if the outer angle in which the circles intersect is large and the opposite label if the angle is small.

In 3D, the algorithm is slightly different since not all Voronoi vertices converge to the medial axis as the sample density approaches infinity. Delaunay tetrahedra called *slivers* are dual to Voronoi vertices that lie arbitrarily close to the surface. They are formed by four points that lie near the equator of a sphere, cf. Figure 2.5. On the other hand, the set of Voronoi vertices called *poles* do converge to the medial axis. The pole of an input site is the point in its Voronoi cell furthest away from the site. This point is always a Voronoi vertex and the empty ball centered at the pole is called the *polar ball*. If we construct the weighted Voronoi diagram from the polar balls and label them as in the two-dimensional case we obtain an appropriate set of approximate maximal balls. The construction of this set of balls is crucial for the construction of input data for the approximation algorithm presented in Chapter 3.

The algorithm in [45] improves the power crust algorithm in the sense that it does not construct a weighted Voronoi diagram of the poles, but constructs an approximate medial axis directly from the Delaunay triangulation of the input sample. To this end, it labels the polar balls using the Delaunay triangulation of the point sample.

2.7 Skin surfaces

The key concept in skin surfaces is the construction of an interpolating radius function. Skin surfaces, introduced by Edelsbrunner in [48], are used for molecular mod-

eling. Molecules are constructed from atoms and each atom is represented by a ball. Smooth patches make the transition between intersecting balls tangent continuous. Edelsbrunner mentions nice properties of skin surfaces:

- *Decomposability*: a skin surface consists of a finite number of degree-2 patches;
- *Constructibility*: there are fast combinatorial algorithms constructing a skin surface or a piecewise linear approximation, see [29, 32] and Chapter 5;
- *Symmetry*: a skin can be defined from the inside as well as the outside;
- *Smoothness*: in the non-degenerate case a skin surface is everywhere tangent continuous;
- *Economy*: even a small number of weighted points can generate skin surfaces with a fairly complicated topological structure;
- *Universality*: every orientable closed surface can be approximated by skin surface;
- *Deformability*: several schemes exist to deform skin surfaces [28, 33, 48], and topological changes of evolving skin surfaces can be efficiently computed;
- *Continuity*: a skin surface varies continuously with points and weights.

2.7.1 Definition

We review the theory of skin surfaces presented in [48]. For further reading on the vector space of weighted points we recommend [84, 41].

In addition to the set of weighted points (balls), a skin surface also depends on an input parameter called the shrink factor. The shrink factor is a scalar $s \in [0, 1]$ that controls the size of the smooth patches over the intersection curves of two balls and also the amount with which the radii of the input balls are shrunk. For a shrink factor is smaller than one, the radii of the input balls become smaller and smooth patches "blend the balls together". For the definition of skin surfaces we use the vector space of the weighted points defined for the construction of the Delaunay triangulation in Section 2.5. Recall the bijective map Π projecting a weighted point $\hat{p}$ to the vector space from Equation (2.2):

$$\Pi(\hat{p}) = (\xi_1, \xi_2, \xi_3, \|p\|^2 - p), \tag{2.2}$$

where $p = (\xi_1, \xi_2, \xi_3)$.

For a shrink factor equal to one, the skin surface of a set of weighted points is obtained by taking the boundary of the union of all weighted points in $\mathrm{conv}(\mathcal{P})$. We show that in that case the skin surface is the boundary of the union of the input balls.

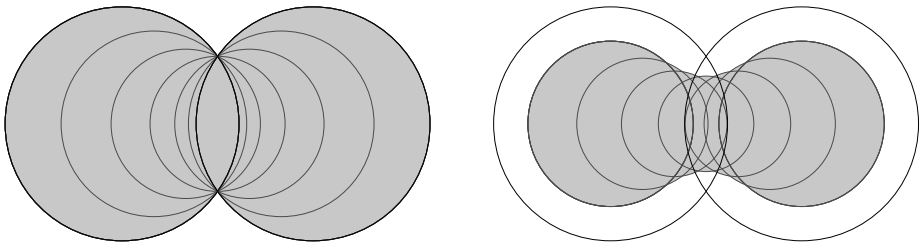

Figure 2.6: The convex hull of two input balls and the same balls shrunk with $s = 0.5$

In the proof of Observation 2, we showed that for a weighted point $\hat{q} = \sum_i \gamma_i \hat{p}_i$, with $\sum_i \gamma_i = 1$ we have:

$$\pi(\hat{q}, x) = \sum_i \gamma_i \, \pi(\hat{p}_i, x) \tag{2.5}$$

Let x be a point outside the union of the input set, then $\pi(\hat{p}, x) > 0$ for all $\hat{p} \in \mathcal{P}$. Hence, the distance of any point $\hat{q} \in \text{conv}(\mathcal{P})$ to x is positive. Therefore all balls in $\text{conv}(\mathcal{P})$ are contained in the union of the input balls and the skin surface for a shrink factor equal to one is the union of the input balls.

The body of a skin surface $\text{bdy}^s(\mathcal{P})$ is the union of the weighted points in $\text{conv}(\mathcal{P})^s$, and the skin surface $\text{skn}^s(\mathcal{P})$ is its boundary:

$$\text{bdy}^s(\mathcal{P}) = \cup(\text{conv}\mathcal{P})^s \tag{2.6}$$

$$\text{skn}^s(\mathcal{P}) = \partial \, \text{bdy}^s(\mathcal{P}). \tag{2.7}$$

Here $\text{conv}(\mathcal{P}) \subset \mathbb{R}^3 \times \mathbb{R}$ is the convex hull – with respect to the vector space structure inherited under Π – of a set of weighted points $\mathcal{P}$, whereas ∂ denotes the boundary – in three space – of the union of the set of balls. Note that this definition generalizes the definition for a shrink factor equal to one.

Figure 2.6 depicts a skin curve in 2D associated with two weighted points. The left figure shows the skin surface with a shrink factor equal to one. Note that every circle passes through the intersection points of the two input circles. On the right, the skin surface for a shrink factor of 0.5 is shown. Since large circles are shrunk more than smaller circles, a smooth transition between the two shrunk input circles appears.

It follows directly from the definition of shrinking that $\text{bdy}^s(\mathcal{P}) \subsetneq \text{bdy}^{s'}(\mathcal{P})$ if $s < s'$. Another important feature of a skin surface is that varying the shrink factor does not change the topology, i.e., for $0 < s, s' \leq 1$, $\text{skn}^s(\mathcal{P})$ and $\text{skn}^{s'}(\mathcal{P})$ are homeomorphic. For $s = 1$, the body of the skin surface is the union of the balls in the input set. In particular, the body of a skin surface is homeomorphic to the union of the balls in the input set.

2.7.2 Relation to the Delaunay triangulation

At first sight, the definition of a skin surface given in Equation (2.7) does not have nice properties. In this section we have a closer look at the weighted points defining

the skin surface. The main result of this section is formulated in the following lemma.

Lemma 4. *For a set of weighted points $\mathcal{P}$ we have*

$$\bigcup \text{conv}(\mathcal{P}) = \bigcup \{\text{conv}(\mathcal{X}) \mid \delta_{\mathcal{X}} \in \text{Del}(\mathcal{P})\}.$$

Proof. Let $\hat{p}$ be a weighted point in $\text{conv}(\mathcal{P})$. There are many weighted points in $\text{conv}(\mathcal{P})$ with the same location p. They only differ in their weight. The weighted point with the largest weight corresponds to the ball with the largest radius. Since their centers are all equal this weighted point contains the other weighted points.

In the vector space structure inherited under Π, the last coordinate of a lifted point decreases if its weight increases. Hence, all weighted points with maximal weight lie on the lower convex hull of the lifted points in $\mathcal{P}$. In Lemma 3 we observed that this lower convex hull is defined by convex combinations of sets of weighted points that define Delaunay simplices.

To conclude, convex combinations of weighted points defining a Delaunay simplex have the property that they maximize weight. Therefore, a these weighted points contain all weighted points in $\mathcal{P}$. □

In view of the this lemma we have:

$$\text{bdy}^s(\mathcal{P}) = \bigcup \{\text{bdy}^s(\mathcal{X}) \mid \delta_{\mathcal{X}} \in \text{Del}(\mathcal{P}).$$

In other words, although a skin surface may have a complicated structure, it can be decomposed into simple pieces defined by at most four weighted points.

2.7.3 Envelopes of weighted points

We analyze the structure of a skin surface defined by up to four weighted points. By definition, $\text{bdy}^s(\mathcal{X}) = \cup\text{conv}(\mathcal{X})^s$. For now, we are not interested in the boundary of $\text{conv}(\mathcal{X})$ and we consider the boundary of $\cup\text{aff}(\mathcal{X})^s$.

If $|\mathcal{X}| = 1$, then $\text{aff}(\mathcal{X})$ contains a single weighted point, say $\hat{p}$ and $\cup\text{aff}(\mathcal{X})^s = \hat{p}^s$. In words, the skin surface of a single weighted point $\hat{p}$ is the sphere bounding the shrunk weighted point $\hat{p}^s$.

In general, the set $\text{aff}(\mathcal{X})$ is an ℓ-dimensional affine subspace of the vector space inherited under Π, with $\ell = |\mathcal{X}| - 1$. By Observation 2, every weighted point in $\text{aff}(\mathcal{X})$ is orthogonal to the orthosphere $\hat{q}$ of $\mathcal{X}$. Assume that $\hat{q}$ is centered at the origin and the affine space is spanned by the first ℓ coordinate-axes. We parameterize $\text{aff}(\mathcal{X})$ as follows. Let $\hat{p}(t_1, \ldots, t_\ell)$ be the weighted point with center $p(t_1, \ldots, t_\ell) = (t_1, \ldots, t_\ell, 0, \ldots, 0)$ and orthogonal to $\hat{q}$. Since $\pi(\hat{p}, \hat{q}) = 0$, it follows that the weight of $\hat{p}(t_1, \ldots, t_\ell)$ is:

$$P(t_1, \ldots, t_\ell) = \| p(t_1, \ldots, t_\ell) \|^2 - Q.$$

Hence, the skin surface is the envelope of the function f defined by:

$$f(t_1, \ldots, t_\ell, x) = \pi(\hat{p}^s(t_1, \ldots, t_\ell), x).$$

A calculation similar to the example in Section 2.3, shows that the envelope of these balls is implicitly defined the zero set of

$$f(\xi_1, \xi_2, \xi_3) = \sum_{i=1}^{\ell} -\frac{1}{1-s}\xi_i^2 + \sum_{j=\ell+1}^{3} \frac{1}{s}\xi_j^2 - Q,$$

where Q is the weight of the orthosphere of $\mathcal{X}$.

Concluding, the envelope of $\mathrm{aff}(\mathcal{X})^s$ is a quadric. Let X be the set of centers of the weighted point in $\mathcal{X}$. For $\ell = 0$, this quadric is a sphere, which is the shrunk weighted point in $\mathcal{X}$. For $\ell = 1$, the envelope is a hyperboloid with symmetry axis $\mathrm{aff}(X)$ and for $\ell = 2$, it is another hyperboloid with symmetry plane $\mathrm{aff}(X)$. For $\ell = 3$, the envelope is a sphere, but now the body of the skin surface lies outside the sphere.

The next step is to analyze how these envelopes contribute to the skin surface. This results in the construction of the mixed complex as we will see in the next section.

2.7.4 Mixed complex

In this section we define a polyhedral structure called the mixed complex $\mathrm{Mix}^s(\mathcal{P})$ that partitions $\mathbb{R}^3$ into convex polyhedra. This partitioning has the property that the mixed cell $\mu^s_{\mathcal{X}}$ contains the part of the skin surface that is defined by the weighted points in $\mathrm{conv}(\mathcal{X})$:

$$\mu^s_{\mathcal{X}} \cap \mathrm{skn}^s(\mathcal{P}) = \mu^s_{\mathcal{X}} \cap \mathrm{skn}^s(\mathcal{X}). \tag{2.8}$$

The mixed complex $\mathrm{Mix}^s(\mathcal{P})$, associated with a shrink factor $s \in [0, 1]$, is an intermediate complex between the Delaunay triangulation and the Voronoi diagram. The Minkowski sum of two sets A and B is the pointwise sum of the points in A and B:

$$A \oplus B = \{a + b \mid a \in A, b \in B\}$$

Each mixed cell in the mixed complex is obtained by taking Minkowski sums of shrunk Delaunay simplices and their dual Voronoi cells.

Definition 5. *For $\delta_{\mathcal{X}} \in \mathrm{Del}(\mathcal{P})$ the* mixed cell $\mu^s_{\mathcal{X}}$ *is defined by* $\mu^s_{\mathcal{X}} = (1-s)\cdot\delta_{\mathcal{X}} \oplus s\cdot\nu_{\mathcal{X}}$.

Here $\cdot$ denotes the multiplication of a set by a scalar and $\oplus$ denotes the Minkowski sum. For $s = 0$ the mixed cell is the Delaunay cell. If s increases it deforms affinely into the Voronoi cell for $s = 1$.

Each mixed cell is a convex polyhedron since it is the Minkowski sum of two convex polyhedra. In 3D, there are four types of mixed cells based on the dimension of the Delaunay simplex. A mixed cell of type ℓ corresponds to a Delaunay ℓ-cell and is of the form $\mu^s_{\mathcal{X}}$ with $|\mathcal{X}| = \ell + 1$. In 3D, mixed cells of type 3 are tetrahedra (shrunk Delaunay 3-cells) and mixed cells of type 0 are shrunk Voronoi 3-cells. A mixed cell of type 1 or 2 is a prism with respectively the shrunk Voronoi facet or the shrunk Delaunay facet as its base.

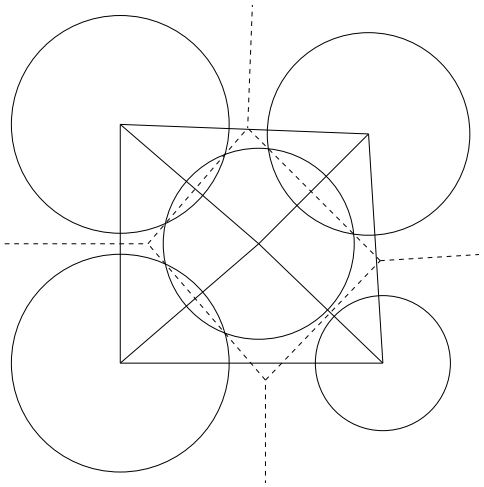

(a) The Voronoi diagram (dashed) and Delaunay triangulation (solid) of five weighted points

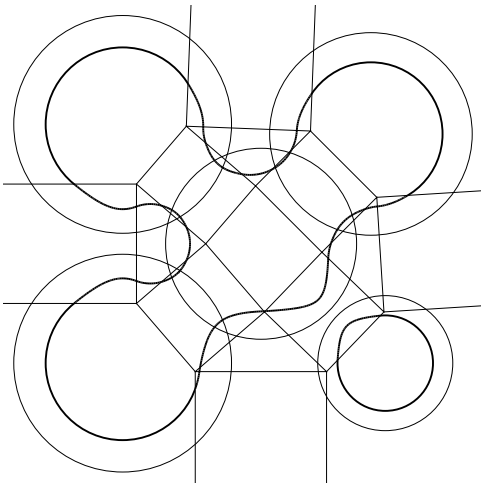

(b) The mixed complex

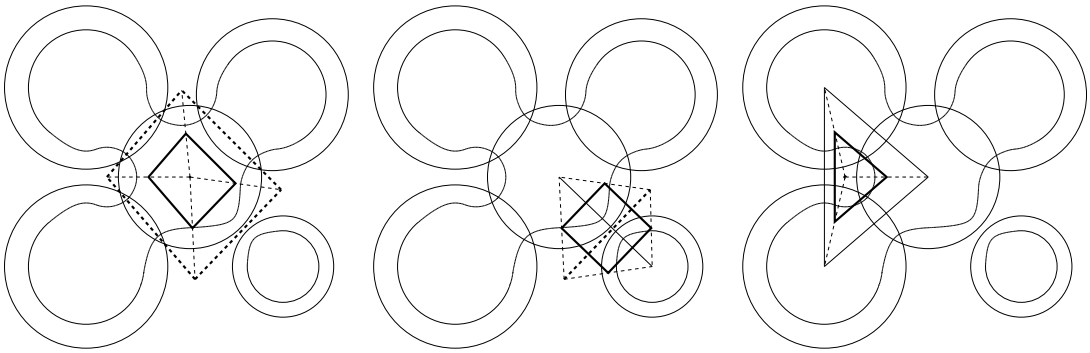

(c) Three types of mixed cells $\mu^s_{\mathcal{X}}$, with $\delta_{\mathcal{X}}$ a vertex, an edge and a face

Figure 2.7: The decomposition of a skin curve of five weighted points.

The intersection of the skin surface and a mixed cell is a piece of a sphere or a hyperboloid see [48]. In fact the mixed cell $\mu^{\mathrm{s}}_{\mathcal{X}}$ defines the region where $\mathrm{skn}^{\mathrm{s}}(\mathcal{X})$ touches the skin surface $\mathrm{skn}^{\mathrm{s}}(\mathcal{P})$, cf. Equation (2.8).

In the plane, the intersection of a skin curve with a mixed cell is either part of a circle or hyperbola. An example of the mixed complex and a skin curve is given in Figure 2.7. In 2D, there are three types of Delaunay simplices and therefore also three types of mixed cells. All rectangles are mixed cells of type 1 and contain hyperbolic patches. The other cells contain circular arcs. Depending on whether the mixed cell is of type 0 or 2, the interior of the skin curve lies inside or outside the circle.

Chapter 3

Approximation by Skin Surfaces

This chapter appeared as [72, 74].

In this chapter, we present a method to approximate a simple, regular C^2 surface W in $\mathbb{R}^3$ by a (tangent continuous) skin surface S. The input of our algorithm is a set balls, the boundary of which is homeomorphic to W and has a small Hausdorff distance to W. By generating patches of spheres and hyperboloids over the intersection curves of the balls the algorithm determines a one-parameter family of skin surfaces, where a parameter controls the size of the patches. The skin surface S is homeomorphic to W and the approximate maximal balls of W in the input set are also maximal with respect to S. See Section 2.6 for the definitions. The Hausdorff distance between the regions enclosed by the input surface W and the approximating skin surface S depends linearly on the Hausdorff distance between the boundary of the union of the input balls and W.

3.1 Introduction

We consider the problem of approximating a simple, regular smooth (C^2) closed surface in $\mathbb{R}^3$ by a skin surface. An example in 2D of a curve, reconstructed using this approach, is drawn in Figure 3.1c. In Figure 3.2 some skin surfaces are drawn for different values of the shrink factor.

The input for the approximation algorithm is a finite sample of the medial axis transform of a surface W, i.e., a finite set of W-maximal balls. Let Σ be a closed surface in $\mathbb{R}^3$ (compact without boundary), which is the boundary of a compact subset V of $\mathbb{R}^3$. A ball is called an *empty ball* of Σ if it is contained in the closure of V. A Σ-*maximal ball* is an empty ball of Σ not contained in any other empty ball of Σ. The *medial axis transform* of W is the set of W-maximal balls. The *medial axis* M of W is the closure of the set of centers of the maximal balls of W, and is a skeletal structure of the surface. See also Section 2.6.

For a sufficiently dense finite subset of the medial axis transform, the boundary of the union of the corresponding maximal balls, forms a good approximation of the surface. Figure 3.1b illustrates this observation for a curve in the plane. Obviously, this approximation is not tangent continuous. Our algorithm reconstructs a tangent

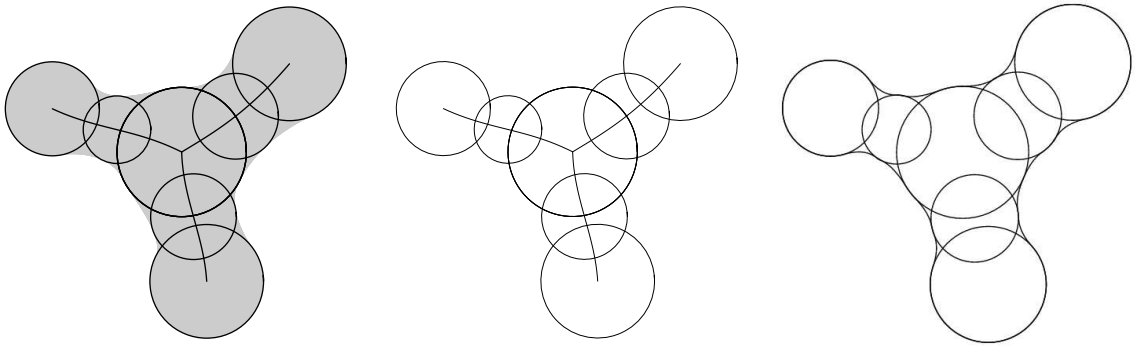

(a) (b) (c)

Figure 3.1: The approximation of a curve in the plane. (a) The medial axis and some maximal circles. (b) A finite set of maximal circles is the input of our algorithm. (c) The output is a skin curve approximating the input curve in Figure 3.1a.

continuous surface from a sample of (approximate) maximal balls by adding smooth patches over the points of intersection. By controlling the size of these patches, we guarantee that the approximation is homeomorphic to the union of the maximal balls. Furthermore, as we increase the sampling density of the set of (approximate) maximal balls, the maximal distance between the regions enclosed by W and the approximating skin surface tends to zero.

We give a more precise specification of the algorithm in Section 3.2, and state the main result on its complexity: the approximating skin surface can be computed in $O(N^2 \log N)$ time and $O(N^2)$ space, where N is the number of approximate maximal balls in the input set. A key ingredient of the algorithm is the construction of the weighted Voronoi diagram of a set of balls with radii growing with the same multiplicative factor. In Section 3.3 we present an algorithm that maintains this diagram as the radii are growing. In Section 3.4 we show how the algorithm satisfies the constraint that the balls in the input set are also S-maximal. Section 3.5 derives a bound on the Hausdorff distance between the regions enclosed by the input surface W and the skin surface S in terms of a parameter related to the sampling density of the set of (approximate) maximal balls. For the proofs in Section 3.3, 3.4 and 3.5 we impose some further constraints on the set of maximal balls. In Section 3.6 we show that the Power Crust algorithm [6] can be used to construct a sample of maximal balls satisfying these constraints.

Related work Amenta et al. [6] show that the medial axis transform of a surface can be effectively approximated from a sufficiently dense sample of points on the surface. If the sample is sufficiently dense, the boundary of the set of approximating maximal balls is homeomorphic to the original surface. In [6] Amenta et al. use the approximation of the medial axis transform to compute the power crust, which approximates the original surface by piecewise linear patches.

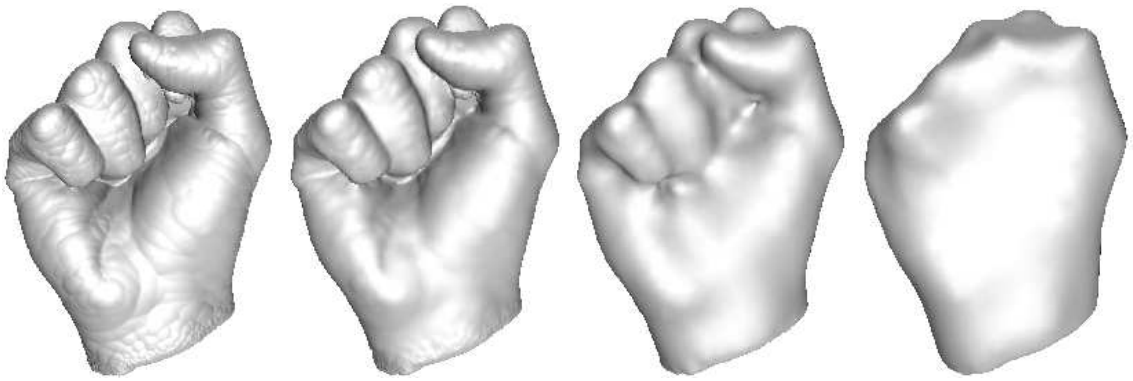

Figure 3.2: The approximating skin surface of a set of balls forming a hand. The patches between the maximal balls become larger as the shrink factor is decreased. The shrink factors (from left to right) are 0.95, 0.8, 0.5 and 0.2.

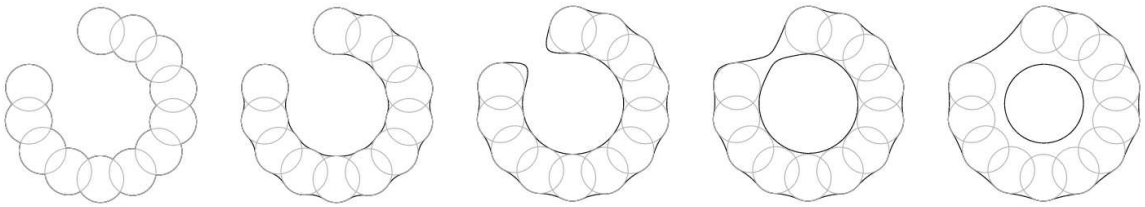

(a) (b) (c) (d) (e)

Figure 3.3: The set of maximal circles and the approximating skin curve for different shrink factors (decreasing from left to right).

Recently, another approach is proposed to model a surface with a skin surface [31]. In this method the radii of the input balls are shrunk during the construction of the approximating skin surface. Our algorithm deliberately prevents shrinking.

3.2 Approximation algorithm

In this section we present the conditions imposed on the surface computed by the approximation algorithm. A maximal ball is a ball contained in a surface W that is maximal with respect to inclusion. We say that a ball $\hat{p}$ is W-maximal if it is a maximal ball of W. For a definition of the Hausdorff distance we refer to Section 2.2.

Definition 6. *Let W be a C^2-surface embedded in three space. For $\varepsilon > 0$ a finite set $\mathcal{P}$ of balls with union M is ε-admissible if*

(i) the boundary ∂M of M is homeomorphic to W;

(ii) every ball in $\mathcal{P}$ is M-maximal;

(iii) The Hausdorff distance between M and the body enclosed by W is at most ε.

For a given C^2-surface W and a positive constant ε, the approximation algorithm takes an ε'-admissible set $\mathcal{P}$ of balls, $0 < \varepsilon' < \varepsilon$, and computes a C^1 *skin surface* S associated with $\mathcal{P}$ satisfying the following properties:

$C_1(S)$**:** W and S are homeomorphic;

$C_2(S)$**:** Every ball in $\mathcal{P}$ is S-maximal;

$C_3(S)$**:** The Hausdorff distance between the bodies of W and S does not exceed ε.

The conditions for an ε-admissible set are rather strong, but not all conditions are needed for the properties to hold. In fact, we only need condition (i) for $C_1(S)$, (ii) for $C_2(S)$ and (iii) for $C_3(S)$. In Section 3.6 we show that the *power crust algorithm* [6] can be used to generate an ε-admissible set.

From the definition of shrinking, Equation (2.3), it follows that $(\hat{p}^{1/s})^s = \hat{p}$. Therefore, $\mathrm{skn}^s((\mathcal{P}^{1/s}))$ contains the balls in $\mathcal{P}$. We define our approximating skin surface $S_s(\mathcal{P})$ as follows.

Definition 7. *For a set of balls $\mathcal{P}$ the surface $S_s(\mathcal{P})$, with $0 < s \leq 1$, is the skin surface with shrink factor s associated with the set of balls $\mathcal{P}^{1/s}$.*

If no confusion is possible, we write S_s instead of $S_s(\mathcal{P})$. If $\mathcal{P}$ is an ε-admissible set of balls then all three properties $C_1(S_s)$–$C_3(S_s)$ are satisfied. Indeed, $\mathcal{P}^1 = \mathcal{P}$, so S_1 is the boundary of the union of these balls and the properties are trivially satisfied. Obviously, S_1 is *not smooth*. However, also for values of s slightly smaller than 1 the properties are satisfied, and S_s is a C^1-surface. Indeed, for s sufficiently near 1, the surface S_s is homeomorphic to S_1, and hence to W. In general each ball touches the union of the balls in a spherical patch. For s near 1, the patches smoothly connecting the input balls are small and therefore all input balls contribute to the skin surface in a spherical patch. These patches make the balls maximal. Because the patches are small, the Hausdorff distance between these patches and the union of the balls is also small.

One or more of these conditions may be violated for smaller values of s. To illustrate these violations in the 2D-case we depict a sequence of skin curves corresponding to decreasing values of s in Figure 3.3. Figure 3.3a shows the skin curve for $s = 1$. As we decrease the shrink factor, the skin curve becomes tangent continuous, due to the appearance of hyperbolic patches connecting the circles, see Figure 3.3b.

Decreasing the shrink factor, i.e., growing these patches, causes a change in the topology of the skin curve, i.e., a violation of condition $C_1(S_s)$, see Figure 3.3d. Finally, as we decrease the shrink factor even further, the balls are no longer maximal, i.e., a violation of $C_2(S_s)$, as depicted in Figure 3.3e. In Section 3.3, we show how to compute the shrink factor at which the change of topology occurs. In Section 3.4 we derive, how far the shrink factor can be decreased without loss of maximality of the balls.

Our goal is to determine the interval of s-values for which conditions $C_1(S_s)$–$C_3(S_s)$ are satisfied. One of the main results of this chapter is

Theorem 8. *A value* $s_i(\mathcal{P}) \leq 1$ *such that for* s, *with* $s_i(\mathcal{P}) < s < 1$, *condition* $C_i(S_s)$ *is satisfied for* $i = 1, 2, 3$, *can be computed in* $O(n^2 \log n)$ *time and* $O(n^2)$ *space, where* n *is the number of balls in* $\mathcal{P}$.

In fact the computed value $s_1(\mathcal{P})$ is the minimal value for which the theorem holds. The algorithm presented in Section 3.3 computes $s_1(P)$ by maintaining the weighted Voronoi diagram of a set of balls with growing radii; This algorithm is of some independent interest. In Section 3.4 this method is slightly adapted to compute an interval of s-values for which condition $C_2(S_s)$ holds. Finally, in Section 3.5, we determine a bound on the error with respect to the Hausdorff distance between the input surface and the computed skin surface. This error analysis gives us a value for $s_3(\mathcal{P})$.

3.3 Maintaining the topology of the union of growing balls

In this section we derive properties of a set of balls with growing radii, which are of independent interest. We introduce a growth factor t, and later on we apply the results of this section for $t = 1/s$, where s is the shrink factor.

We compute the smallest t, such that the boundary of the union of the balls in $\mathcal{P}^t$ is not homeomorphic to $\partial \cup \mathcal{P}$, viz. the transition of Figure 3.3c–d. Since the skin surface $S_s(\mathcal{P})$ is homeomorphic to the boundary of the union of the set of balls $\mathcal{P}^t$, see Section 2.7, this is equivalent to computing the value $s_1(\mathcal{P})$, defined in Section 3.2. Figure 3.4 illustrates this equivalence in 2D, where the change of topology of the skin curve coincides with the change of topology of the union of balls defining the skin curve.

Definition 9. *The boundary* $\partial \cup \mathcal{P}^t$ *of the union of balls in* $\mathcal{P}^t$ ***changes topology*** *at* $(x, \tau) \in \mathbb{R}^3 \times \mathbb{R}$ *if for every neighborhood* U *of* x *in* $\mathbb{R}^3$ *there is an* $\varepsilon > 0$ *such that, for* $\tau - \varepsilon < t_1 < \tau < t_2 < \tau + \varepsilon$, *the sets* $U \cap \partial \cup \mathcal{P}^{t_1}$ *and* $U \cap \partial \cup \mathcal{P}^{t_2}$ *are not homeomorphic.*
A ball $\hat{q} \in \mathcal{P}$ *is* ***involved*** *in the change of topology at* (x, τ) *if* x *lies on the sphere bounding* $\hat{q}^\tau$.

Obviously, if $\partial \cup \mathcal{P}^t$ changes topology at (x, τ), then $x \in \partial \cup \mathcal{P}^\tau$. Occasionally, we just say that $\partial \cup \mathcal{P}^t$ changes topology for $t = \tau$ if the location x is irrelevant.

The following result states that a set of weighted points changes topology if their intersection becomes non-empty.

Lemma 10. *If all weighted points of* $\mathcal{Q}$ *are involved in a change of topology of* $\partial \cup \mathcal{Q}^t$ *at* (x, τ), *then*

$$\cap \mathcal{Q}^t = \begin{cases} \emptyset, & \text{for } t < \tau, \\ \{x\}, & \text{for } t = \tau, \end{cases}$$

and $\cap \mathcal{Q}^t \neq \emptyset$, *for* $t \geq \tau$.

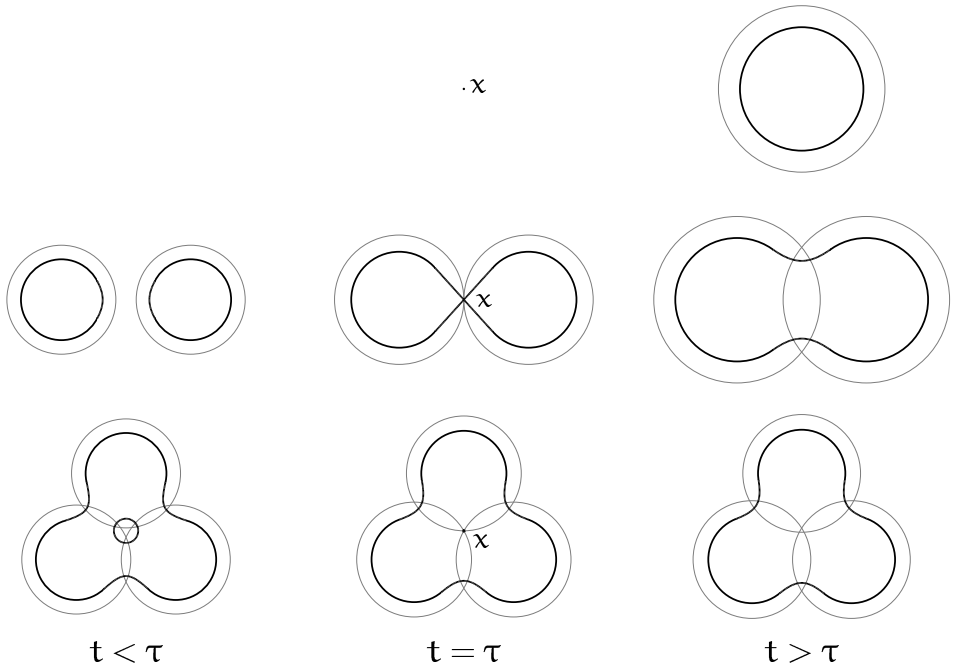

Figure 3.4: Change of topology at (x, τ) of a union of one, two and three weighted points.

Proof. Since all weighted points in $\mathcal{Q}$ are involved, the intersection of $\mathcal{Q}^\tau$ is not empty. For $t_0 > \tau$, the ball $\hat{p}^\tau$ is contained in $\hat{p}^{t_0}$, hence $\cap\mathcal{Q}^{t_0} \neq \emptyset$. Because a ball is convex, the union of $\mathcal{Q}^{t_0}$ is star shaped for $t_0 > \tau$. Its boundary is homeomorphic to a 2-sphere and does not change topology for $t_0 > \tau$.

If $\cap\mathcal{Q}^{t_0} \neq \emptyset$, for $t_0 < \tau$, then $\partial \cup \mathcal{Q}^{t_0}$ is homeomorphic to a 2-sphere. Hence, $\partial \cup \mathcal{Q}^{t_0}$ does not change topology for $t_0 = \tau$, which contradicts the prerequisite. □

It follows from the proof that, for $t > \tau$, the union of $\mathcal{Q}^t$ is homeomorphic to a ball. Furthermore, x is an interior point of $\cap\mathcal{Q}^t$.

Lemma 11. *If $\partial \cup \mathcal{P}^t$ changes topology at $(x, \tau) \in \mathbb{R}^3 \times \mathbb{R}$, then the subset $\mathcal{Q}$ of $\mathcal{P}$ of weighted points involved in this change of topology defines a cell in the weighted Voronoi diagram of $\mathcal{P}^\tau$. This cell contains the point x.*

Proof. It follows from Lemma 10 that x belongs to the boundary of the weighted points in $\mathcal{Q}^\tau$, so $\pi(\hat{q}^\tau, x) = 0$, for $\hat{q} \in \mathcal{Q}$. for $\hat{p} \in \mathcal{P} \setminus \mathcal{Q}$, the distance $\pi(\hat{p}^\tau, x) > 0$, since $x \notin \hat{p}^\tau$. Hence, x belongs to the cell $\nu_{\mathcal{Q}^\tau}$ of the weighted Voronoi diagram of $\mathcal{P}^\tau$. □

Consider a set $\mathcal{Q}$ of balls, such that the boundaries of the weighted points intersect in x. Generically, the tangent planes at x intersect transversally, i.e., the intersection is an affine space of co-dimension $|\mathcal{Q}|-1$. The balls $\mathcal{Q}^t$ form a one-parameter family, so at isolated values of t we expect intersections with non-transversal tangent planes, i.e., the co-dimension of this intersection *is less than* $|\mathcal{Q}|-1$. We impose the following *generic condition* on the family $\mathcal{P}$.

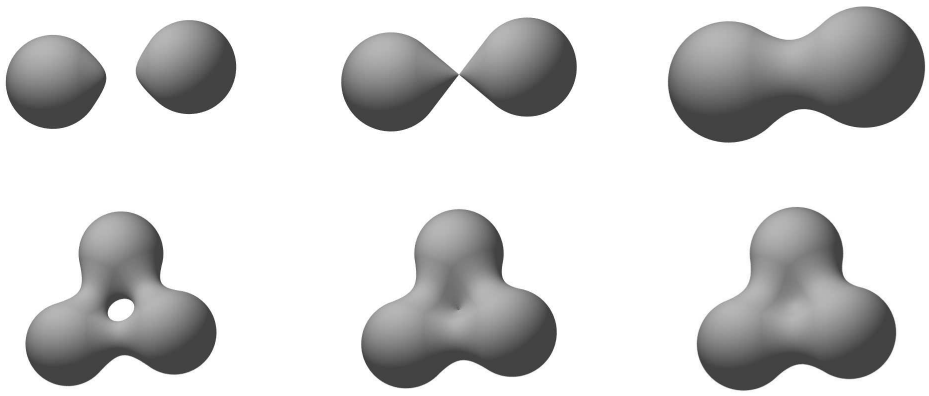

Figure 3.5: The topological changes of a skin surface defined by two and three weighted points as the input balls are grown.

Generic change of topology. *If $\partial \cup \mathcal{P}^t$ changes topology at (x, τ), and the subset $\mathcal{Q}$ of $\mathcal{P}$ consists of the balls involved in this change of topology, then the tangent planes at x of the balls in $\mathcal{Q}^\tau$ intersect in an affine space of co-dimension $|\mathcal{Q}| - 1$.*

In the plane, at most three weighted points are involved in a generic change of topology, see Figure 3.4. A single weighted point $\hat{p}$ becomes real and creates a new component at $t = 0$. If two weighted points are involved, the change in topology corresponds to the *creation of a bridge* between two parts of the boundary. A change of topology in which three weighted points are involved corresponds to the *filling of a void*.

To describe the situation in three-space, let $\mathcal{Q}$ be the set of balls involved in a generic change of topology, then $|\mathcal{Q}| \leq 4$. The changes of topology correspond to the creation of a new component, if $|\mathcal{Q}| = 1$, the creation of a bridge, if $|\mathcal{Q}| = 2$, the filling of a tunnel, if $|\mathcal{Q}| = 3$, and the filling of a void, if $|\mathcal{Q}| = 4$.

Lifting the weighted Voronoi diagram In view of Lemma 11, topological changes in $\partial\mathcal{P}^t$ are related to the cells of the weighted Voronoi diagram. To incorporate t-dependence, we lift the weighted Voronoi diagram to the extended space $\mathbb{R}^3 \times \mathbb{R}$. To this end let $\hat{p}, \hat{q} \in \mathcal{P}$ and consider the halfspace $H^*_{\hat{p},\hat{q}} \subset \mathbb{R}^3 \times \mathbb{R}$ defined by

$$\begin{aligned} H^*_{\hat{p},\hat{q}} &= \{(x,t) \in \mathbb{R}^3 \times \mathbb{R} \mid \pi_{\hat{p}^t}(x) \leq \pi_{\hat{q}^t}(x)\} \\ &= \{(x,t) \in \mathbb{R}^3 \times \mathbb{R} \mid \\ &\qquad 2\langle x, q - p\rangle + t(Q - P) \leq \|q\|^2 - \|p\|^2\}. \end{aligned}$$

The boundary of this halfspace is called the *extended bisector* of the weighted points $\hat{p}$ and $\hat{q}$. With a weighted point $\hat{p} \in \mathcal{P}$ we associate the extended Voronoi cell $V^*(\hat{p})$ in $\mathbb{R}^3 \times \mathbb{R}$ defined by

$$V^*(\hat{p}) = \bigcap_{\hat{q}\,:\,\hat{q} \neq \hat{p}} H^*_{\hat{p},\hat{q}}.$$

For $\hat{p} = (p, p) \in \mathcal{P}$, the point $(p, 0)$ belongs to $V^*(\hat{p})$, so no extended Voronoi cell is empty. The extended Voronoi cells determine a subdivision of $\mathbb{R}^3 \times \mathbb{R}$ into convex polyhedra. This subdivision is called the *extended Voronoi diagram* of $\mathcal{P}$, and is denoted by $VD^*(\mathcal{P})$. The weighted Voronoi cell of $\hat{p}^\tau$ with respect to $\mathcal{P}^\tau$ is the intersection of the extended Voronoi cell $V^*(\hat{p})$ and the hyperplane $t = \tau$.

General position of scaled weighted points. *We assume that the weighted points in $\mathcal{P}$ are in general position in the sense that*

1. *The centers of $k+2$ weighted points in $\mathcal{P}$ lie on a common k-flat, for $0 \leq k \leq 2$;*
2. *no 6-tuple of weighted points in $\mathcal{P}^t$ has a common orthosphere;*
3. *a 5-tuple of weighted points in $\mathcal{P}^t$ has a common orthosphere for an unique value of t.*

For a set of weighted points in general position, there is a constant number of 1-faces and 2-faces incident upon a vertex of the extended Voronoi diagram.

If $\mathcal{P}$ is in general position, then, for $0 \leq k \leq 4$, the k-flat supporting a k-face of the extended Voronoi diagram is defined by $5-k$ weighted points, in other words, it is incident to $5-k$ extended Voronoi cells (a cell is a 4-face of the extended Voronoi diagram). Since a vertex of the extended Voronoi diagram is defined by 5 weighted points, it is incident upon $\binom{5}{4} = 5$ edges. The number of edges incident upon a vertex (x, τ) of the extended Voronoi diagram, and contained in the halfspace below (above) the hyperplane $t = \tau$, is at least 1 since the extended Voronoi diagram subdivides $\mathbb{R}^3 \times \mathbb{R}$. If a vertex has $k+1$ incident edges in the halfspace below the hyperplane $t = \tau$, for $0 \leq k \leq 3$, then this vertex is said to be of *type* k. Since no Voronoi cell is empty at $t = 0$, we do not encounter vertices of type 0.

Lemma 12. *The extended Voronoi diagram of a set of N weighted points in general position in $\mathbb{R}^3$ has $O(N^2)$ faces.*

Proof. Obviously, the number of 4-faces is $O(N)$, since there is a one-one correspondence between the set of 4-faces and the set of weighted points. Similarly, the number of 3-faces is $O(N^2)$, since each 3-face is defined by a pair of weighted points.

Vertices (0-faces) of type 0 (type 3) are t-minimal (t-maximal) points of a 4-face, so there are $O(N)$ vertices of type 0 and of type 3. Vertices of type 1 (type 2) are t-minimal (t-maximal) points of a 3-face, so there are $O(N^2)$ vertices of type 1 and of type 2.

Five 1-faces are incident to a vertex. Therefore the number of 1-faces is $O(N^2)$. Each 2-face incident to a vertex is also incident to two 1-faces incident to the vertex. Hence, the number of 2-faces incident to a vertex is at most $\binom{5}{2} = 10$ and the number of 2-faces is $O(N^2)$.

Since we assume general position, the total number of faces of the extended Voronoi diagram is $O(N^2)$. □

To see that the quadratic complexity is attainable, we construct a set of weighted points for which the structure of the weighted Voronoi diagram changes a quadratic

Figure 3.6: A set of weighted points for which the Voronoi diagram changes topology a quadratic number of times is shown in (a). The figures (b)–(d) show the Voronoi diagram for increasing t.

number of times as we grow the weighted points. More specifically, the Voronoi cells of $O(N)$ weighted point become empty which corresponds to vertices of type 3 in the extended 3D Voronoi diagram. Further, $O(N^2)$ flips of a Voronoi edge to a Voronoi facet and vice versa (these are vertices of type 1 or 2) occur.

The construction is planar. First align $N/2$ points along the vertical axis with weight p. The other weighted points are aligned along the positive horizontal axis and have a weight p', with $p' < p$. Further, the weighted point on the horizontal axis closest to the origin is positioned such that its Voronoi cell is incident to all Voronoi cells of the weighted points positioned on the vertical axis, viz. Figure 3.6(a).

As the weights increase, the Voronoi cells of the weighted points positioned on the vertical become larger, whereas the Voronoi cells of the weighted points on the horizontal axis become smaller. As a result for each weighted point on the horizontal axis, we have $O(N)$ edge flips before the Voronoi cell becomes empty. The outmost weighted point forms an exception, since their Voronoi cells do not become empty. This process is depicted in Figure 3.6(b) to (d).

Proposition 13. *The extended Voronoi diagram of a set of* N *weighted points in* $\mathbb{R}^3$ *can be constructed in* $O(N^2 \log N)$ *time.*

Proof. We maintain the weighted Voronoi diagram of $\mathcal{P}^t$ using a sweep-hyperplane algorithm with planes t = constant, starting at $t = 0$. The weighted Voronoi diagram for $t = 0$ can be computed in $O(N^2)$ time, using the algorithm from [68]. The combinatorial structure of $VD(\mathcal{P}^t)$ changes at the t-coordinates of the vertices of the extended Voronoi diagram. There are several types of *events*, depending on the type of the vertex: the type of the event is by definition the type of the vertex. At an *event of type* k, $0 < k \leq 3$, a k-face f (and its incident j-faces, $0 \leq j < k$) is destroyed, and a $3-k$-face f' is created.

Since the weighted points are in general position, face f is a k-simplex just before its destruction. The algorithm maintains a priority queue storing the destruction times of all k-simplices in the current weighted Voronoi diagram. A Voronoi vertex corresponds to a line segment in the extended Voronoi diagram. By definition of the extended Voronoi diagram, a Voronoi vertex parameterized in terms of t, moves

along a line in $\mathbb{R}^3$ at constant speed. At an event, at least two Voronoi vertices coincide and five weighted points have a common orthosphere.

Processing an event of type k boils down to the following update operations:

(i) Update the weighted Voronoi diagram, i.e., remove the destroyed k-simplex and insert the created $3-k$-simplex, thereby updating all incidence relations;

(ii) Update the priority queue by adjusting the destruction times of all simplices, incident to the newly created simplex.

Adding the destruction time of the newly created simplex is not necessary since it has the value of the current event. In each event at least one edge is destroyed, hence the destruction time of a k-facet is equal to the destruction time of one of its edges. Therefore, it is sufficient to compute the destruction time of each new or updated edge. A constant number of faces is involved in a single event. Hence, step (i) takes $O(1)$ time and step (ii) takes $O(\log N)$ time.

In view of Lemma 12 the total number of time stamps is $O(N^2)$. Therefore the time complexity of the algorithm is $O(N^2 \log N)$. □

In practice, there is no need to compute the entire extended Voronoi diagram and our algorithm terminates as soon as one of the properties of our approximation is violated.

Corollary 14. *Let $\mathcal{P}$ be a set of N weighted points in $\mathbb{R}^3$. Then τ_0, defined by*

$$\tau_0 = \min\{\tau > 1 \mid \partial \cup \mathcal{P}^t \text{ changes topology for } t = \tau\},$$

can be computed in $O(N^2 \log N)$ time.

Proof. Consider a face of the extended Voronoi diagram, defined by a subset $\mathcal{Q}$ of $\mathcal{P}$. There is a unique value t for which $\cap \mathcal{Q}^t$ consists of a single point, x say: solve (x, t) from the equation $\pi(\hat{q}^t, x) = 0$, for some, and hence all, $\hat{q} \in \mathcal{Q}$. In this way a unique time stamp $\tau_{\mathcal{Q}}$ is defined for each k-face of the extended Voronoi diagram. The value τ_0 is the minimal time stamp greater than 1. Since there are $O(N^2)$ faces, it can be computed from the extended Voronoi diagram in $O(N^2)$ additional time. □

3.4 Preserving maximality of the balls

The medial axis transform is a representation of a surface (the boundary of the union of all maximal balls is the surface). An ε-admissible set approximates the medial axis transform of a surface W by a set $\mathcal{P}$ of approximate maximal balls of W. The balls are approximate maximal balls since the body of W and $\cup\mathcal{P}$ have small Hausdorff distance and the balls are maximal in balls of $\partial \cup \mathcal{P}$. In this section, we compute a shrink factor $s_2(\mathcal{P})$, such that all balls in $\mathcal{P}$ are maximal in the approximating

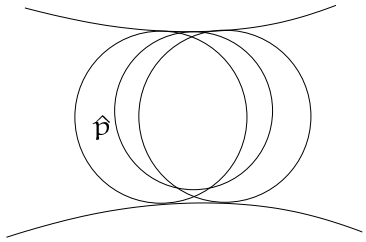

Figure 3.7: Three weighted points in an ε-admissible set $\mathcal{P}$. The radius of the maximal circle $\hat{p}$ in the middle is shrunk with a factor ε. The result is that $\hat{p}$ only touches $\cap\mathcal{P}$ in one small circular arc.

skin surface $S_s(\mathcal{P})$, for $s_2(\mathcal{P}) < s \leq 1$. This means that we do not discard any information given by the maximal balls in $\mathcal{P}$.

Consider an ε-admissible set $\mathcal{P}$. The surface $M = \cup\mathcal{P}$ is composed of spherical patches. Each spherical patch is part of the boundary of some maximal ball in $\mathcal{P}$.

Requirement (ii) of an ε-admissible set ensures that each weighted point generates at least one spherical patch on ∂M, otherwise the weighted point is not ∂M-maximal. In Figure 3.7, we show that for any $r > 0$, we can construct an ε-admissible set that contains a weighted point $\hat{p}$ such that $\hat{p}$ generates only one circular arc and this arc is arbitrarily small. This example can be extended in any dimension. To guarantee that each weighted point in $\mathcal{P}$ is also maximal in the approximating skin surface, we show that each input ball intersects the skin surface in a spherical patch γ.

Conceptually, we start with the skin surface for $s = 1$. Since each ball touches the boundary of the union of the input balls, there exists a spherical patch γ for each weighted point. As we decrease s, smooth patches between spheres arise and $\gamma \cap S_s$ becomes smaller. We stop decreasing s just before the last spherical patch generated by an input ball disappears.

To determine if an input ball intersects the skin surface S_s in a point, we use the mixed complex. Recall that the mixed complex decomposes the skin surface into quadratic patches. In particular, a mixed 0-cell clips its corresponding maximal ball. A mixed 0-cell corresponds to a weighted point $\hat{p} \in \mathcal{P}$. Its shape is the Voronoi 3-cell of $\hat{p}$ shrunk with a factor s towards the weighted Delaunay 0-cell $\delta_{\{\hat{p}\}}$.

Lemma 15. *Let γ be a spherical patch on the boundary of $\cup\mathcal{P}$ generated by a weighted point $\hat{p} \in \mathcal{P}$. Then there is a weighted Delaunay 3-cell $\delta_{\mathcal{X}}$ adjacent to $\delta_{\{\hat{p}\}}$ for which the segment between p and $v_{\mathcal{X}}$ intersects γ. Furthermore, the skin surface S_s contains a spherical patch of γ iff*

$$s \cdot \|p - v_{\mathcal{X}}\|^2 > P,$$

for some 3-cell $\delta_{\mathcal{X}}$ incident to $\delta_{\{\hat{p}\}}$.

Proof. First, we analyze the shape of a mixed 0-cell more carefully. A Voronoi 3-cell is a convex polyhedron with its Voronoi 0-cells at its vertices. Since a mixed 0-cell is a Voronoi 3-cell shrunk towards the center of its corresponding weighted point, a

mixed 0-cell is a convex polyhedron with its vertices on the line from the weighted Delaunay 0-cell to each of the adjacent Voronoi 0-cells.

To be more precise, for every weighted Delaunay 3-cell $\delta_{\mathcal{X}}$ adjacent to $\delta_{\{\hat{p}\}}$, the mixed 0-cell has a vertex at the point:

$$s \cdot \nu_{\mathcal{X}} + (1 - s) \cdot \delta_{\{\hat{p}\}}.$$

The intersection of the spherical cap γ with S_s is non-empty if some part of γ lies inside the mixed 0-cell. Since a mixed 0-cell is a convex polyhedron, this is the case if a vertex of the polyhedron lies outside $\hat{p}$. Therefore the intersection of γ with the skin surface is non-empty if the inequality in the lemma holds. Note that in this equation $\nu_{\mathcal{X}}$ also depends on the shrink factor. □

The algorithm to obtain the lowest shrink factor is similar to the algorithm used to compute the minimal shrink factor for which the skin surface is homeomorphic to the original surface W. Instead of testing if a set of weighted points in a weighted Delaunay cell causes a topological change, we test for each weighted Delaunay vertex if a spherical patch degenerates into a point. Through our study of degeneration, we have obtained the lowest shrink factor for which we can guarantee that the maximal balls in $\mathcal{P}$ are also maximal in the skin surface.

3.5 Error estimates

In this section we prove the following approximation result.

Proposition 16. *Let $\varepsilon > 0$, and let $\mathcal{P}$ be an ε'-admissible sample of a surface W, where $0 < \varepsilon' < \varepsilon$. Then there is an $s_3(\varepsilon, \varepsilon')$, with $0 < s_3(\varepsilon, \varepsilon') < 1$, such that*

1. *$s_3(\varepsilon, \varepsilon')$ can be computed in $O(N^2 \log N)$ time;*
2. *For $s_3(\varepsilon, \varepsilon') \leq s \leq 1$, the Hausdorff-distance between the body B_W bounded by W and the body $B_S(s)$ bounded by $S_s(\mathcal{P})$ does not exceed ε.*

Since $\mathcal{P}$ is ε'-admissible it follows that $d(B_W, \cup\mathcal{P}) \leq \varepsilon'$, see Section 3.2, condition (iii). Therefore, the main result would be clear if we can prove that

$$d(\cup\mathcal{P}, B_S(s)) \leq \varepsilon - \varepsilon'$$

for shrink factors $s_3(\varepsilon, \varepsilon') \leq s < 1$. The latter statement follows from the following result by taking $s_3(\varepsilon, \varepsilon') = s'_3(\varepsilon - \varepsilon')$:

Lemma 17. *For $\eta > 0$ there is an $s'_3(\eta)$, with $0 < s'_3(\eta) < 1$, such that*

1. *$s'_3(\eta)$ can be computed in $O(N^2 \log N)$ time;*
2. *For $s'_3(\eta) \leq s \leq 1$, the Hausdorff-distance between $\cup\mathcal{P}$ and $B_S(s)$ does not exceed η.*

To prove the lemma, we need three technical lemmas. The first result states that, for an efficiently computable range of shrink factors, the body $\mathcal{B}_{\mathcal{P}}(s) = \mathrm{bdy}^s(\mathcal{P}^{1/s})$ of the skin surface $S_s(\mathcal{P})$ intersects a cell $\mu^s_{\mathcal{X}}$ of the mixed complex of $\mathcal{P}^{1/s}$ only if the focus $f(\mathcal{X})$ of $\mathcal{X}$ is contained in $\cup\mathcal{X}$.

Lemma 18. *There is an s''_3, with $0 < s''_3 < 1$, such that*

1. *s''_3 can be computed in $O(N^2 \log N)$ time;*
2. *If $s''_3 \leq s \leq 1$ and $\mu^s_{\mathcal{X}} \cap \mathcal{B}_{\mathcal{X}}(s) \neq \emptyset$, then $f(\mathcal{X}) \in \cup\mathcal{X}$.*

Proof. The main idea of the proof is again to use the sweep-line algorithm and to compute the maximal shrink factor for which $\mu^s_{\mathcal{X}} \cap \mathcal{B}_{\mathcal{X}}(s) = \emptyset$. By taking the maximum of all these local values, we obtain a global maximum. This global maximum is the minimal value for which condition *2* is not satisfied.

In the remainder of this proof we show how to compute the shrink factor locally for different types of mixed cells. The focus of a mixed 0-cell is the center of the generating weighted point, which lies inside the weighted point.

A mixed 1-cell corresponding to a set $\mathcal{X}$ contains a two sheeted hyperboloid. Its symmetry axis is the line through the centers of the two weighted points in $\mathcal{X}$. For $s = 1$, we know that the mixed 1-cell does not contain any part of the skin surface, since no patches are generated. As we decrease the shrink factor, the first point where the skin surface intersects the mixed 1-cell is on the symmetry axis.

A similar analysis of mixed 2-cells and 3-cells, shows that the first point where the skin surface will intersect the mixed cell lies on the segment between the focus $f(\mathcal{X})$ and the center of a weighted point in $\mathcal{X}$. Moreover, the skin surface intersects the mixed cell if

$$s \cdot \| p - f(\mathcal{X}^{1/s}) \|^2 \leq P,$$

for some $\hat{p} \in \mathcal{X}$. This can be shown in a similar way as done in the proof of Lemma 15 and is left as an exercise to the reader. □

As a corollary, for shrink factors in the range $[s''_3, 1]$ the patches of the skin surface are defined by subsets $\mathcal{X}$ of $\mathcal{P}$ for which all weighted points in $\mathrm{aff}(\mathcal{X})$ have positive weights. This is made more precise in our second technical lemma.

Lemma 19. *If $f(\mathcal{X}) \in \cup\mathcal{X}$, then the weight P of a weighted point $\hat{p} \in \mathrm{aff}(\mathcal{X})$ is non-negative.*

Proof. All weighted points in $\mathrm{aff}(\mathcal{X})$ are orthogonal to the orthosphere $\hat{q}$ of $\mathcal{X}$, cf. Observation 2.

A simple, but tedious, calculation shows that for $\hat{p} \in \mathrm{aff}(\mathcal{X})$, $\pi(\hat{p}, f(\mathcal{X})) =$ constant. Since $f(\mathcal{X})$ lies inside $\hat{p}' \in \mathcal{X}$, $\pi(\hat{p}', f(\mathcal{X})) \leq 0$, which means that $\pi(\hat{p}, f(\mathcal{X})) \leq 0$. As a result, $\| p - f(\mathcal{X}) \|^2 \leq P$, so $P \geq 0$. Therefore, all weighted points in $\mathrm{aff}(\mathcal{X})$ have positive weight. □

Our third technical result provides us with an upper bound on the Hausdorff-distance between a quadratic patch and (a subset of) the union of balls in the input

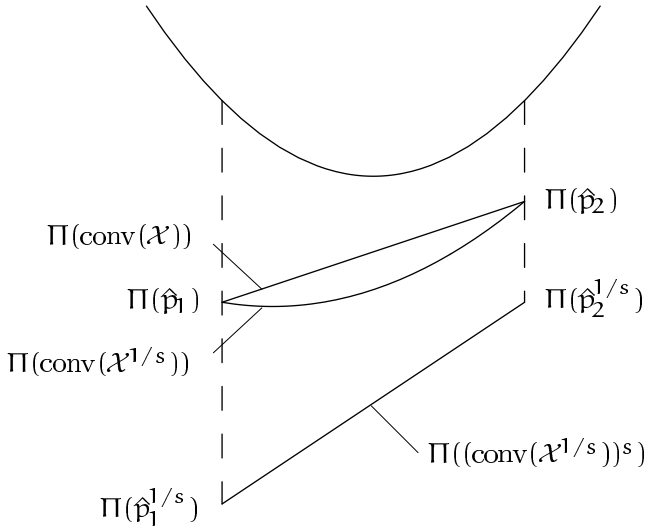

Figure 3.8: The vector space obtained under the lifting operator Π, defined in Equation (2.2), on 1-dimensional weighted points. It shows how the convex hull of the set $\mathcal{X} = \{\hat{p}_1, \hat{p}_2\}$ deforms as we decrease the shrink factor.

set. Let $\mathcal{X}$ be a subset of $\mathcal{P}$ with $f(\mathcal{X}) \in \cup\mathcal{X}$. Let $X = \{p \mid \hat{p} \in \mathcal{X}\}$ and let $m \in \mathrm{aff}(X)$ be the point equidistant to each $x \in X$. Let $d_{\mathcal{X}}$ be the corresponding distance $\|m - x\|$.

Lemma 20. *For* $s_{\mathcal{X}}(\varepsilon) \leq s \leq 1$ *with*

$$s_{\mathcal{X}}(\varepsilon) = 1 - \frac{\varepsilon^2}{d_{\mathcal{X}}^2},$$

we have that

$$d(\cup\mathcal{X}, \mathcal{B}_{\mathcal{X}}(s)) \leq \varepsilon.$$

Proof. We define the function $P_s(p)$ as the weight of the weighted point centered at p in $(\mathrm{conv}(\mathcal{X}^{1/s}))^s$.

The set $\Pi(\mathrm{conv}(\mathcal{X}))$ forms a hyperplane in $\mathbb{R}^4$. Since the weight is the distance along the last coordinate axis to the unit paraboloid, $P_1(p)$ is a paraboloid with leading coefficient 1. For the analogue of the vector space obtained under Π of weighted points with an 1-dimensional center, see Figure 3.8. The top parabola denotes the set of weighted points with zero weight.

A similar reasoning shows that $\Pi(\mathrm{conv}(\mathcal{X}^{1/s}))$ is also a hyperplane in $\mathbb{R}^4$. By the definition of shrinking, $P_s(p)$ is a paraboloid with leading coefficient s.

Knowing the leading coefficient of $P_s(p)$, we now determine its top. Consider a weighted point $\hat{p}' \in \mathcal{X}$. It is first shrunk with a factor $1/s$ and then with a factor s, therefore $P_s(p') = P'$, as is also depicted in Figure 3.8. Using these points the parabola is uniquely defined, and it follows that

$$P_s(p) = P_1(p) - (1 - s)(\|x - m\|^2 - \|p - m\|^2),$$

for some $x \in X$, since $\|x-m\| = \|x'-m\|$ for $x, x' \in X$. Indeed, $P_s(p)$ is a paraboloid with leading coefficient s, since $P_1(p)$ has leading coefficient 1. Further, the weight of the points in $\mathcal{X}$ is independent of the shrink factor.

Using the fact that for $s = 1$, $\hat{p}(1) \in \text{conv}(\mathcal{X})$ is contained in $\cup\mathcal{X}$ and that each weighted point corresponds to a ball with real radius, we bound the patch in terms of the distance between $\hat{p}(s) \in \text{conv}(\mathcal{X}^{1/s})^s$ and $\hat{p}(1)$.

$$\begin{aligned} d(\hat{p}(s), \cup\mathcal{P}) &\leq d(\hat{p}(s), \hat{p}(1)) \\ &= \sqrt{P(s)} - \sqrt{P(1)} \\ &\leq \sqrt{(1-s)(\|x-m\|^2 - \|p-m\|^2)} \end{aligned}$$

The equation in the lemma follows if we take $p = m$ to obtain the maximal difference in weight, and then solve $d(\hat{p}(s), \cup\mathcal{P}) = \varepsilon$ for s. Note that this bound can be too pessimistic since m can lie outside the convex hull of X, and therefore $\|p-m\|^2 > 0$ for all $p \in \text{conv}(X)$ □

Proof of Lemma 17. Since $\cup\mathcal{P} \subset B_S(s)$, for $0 < s \leq 1$, it follows that $d(x, B_S(s)) = 0$ for $x \in \cup\mathcal{P}$. Conversely, let $x \in B_S(s)$. Let $s'_3(\eta)$ be the maximum of s''_3 and

$$\max\{s_{\mathcal{X}}(\eta) \mid f(\mathcal{X}) \in \cup\mathcal{X}\},$$

then the statement follows from Lemmas 18, 19 and 20. □

3.6 Generating input data

This section presents a way to obtain an ε-admissible set of maximal balls from a finite sample of points on the surface W bounding a volume B_W. In the computational geometry community several algorithms [3, 6, 43] have been proposed to reconstruct a surface from a finite point sample on the surface. Some of these algorithms first approximate the medial axis (transform) and use it for reconstruction. In this section, we show that the approximate medial axis transform of the power crust [6] can also be used to generate an ε-admissible sample of maximal balls.

The extended medial axis transform is the medial axis transform not restricted to the inside of the surface. All maximal balls in the extended medial axis transform touch the surface in at least two points, and contain no surface points in their interior. The extended medial axis is the set of centers of the maximal balls in the extended medial axis transform.

The power crust We will only briefly discuss the power crust here and refer to [6] for a more detailed discussion. The method is based on the observation that the Voronoi cells of a dense sample X on a surface W are elongated in the direction of the normals.

To be more precise, the *poles* of a site $x \in X$ are the two points in the Voronoi cell of x farthest away from x, such that the two points lie on different sides of W. The

balls centered at the poles that touch the sites are called the *polar balls*. The main observation of Amenta and Bern [3] is that the two poles lie close to the extended medial axis, one to the interior and the other to the exterior medial axis.

If the polar balls are found, it remains to label the balls in- and outside. Therefore, consider the Voronoi diagram of the polar balls. Start with labeling one polar ball on the convex hull as outside and label adjacent polar balls according to the following three criteria: first, two polar balls corresponding to the same site have opposite labels. Secondly, if the two balls have a shallow intersection, they lie on opposite sides of the surface, because of the assumption that distance between sample points is small with respect to the radius of the polar balls. On the other hand, if two balls intersect deeply, they have the same label.

To complete the surface reconstruction process, the Voronoi facets separating inner from outer poles are returned.

Sampling of the surface Under some conditions on the point sample of the surface, the power crust can be proven to give accurate results. The following definitions are taken from [3, 4].

Definition 21. *The* Local Feature Size *(*LFS*) at a point $w \in W$ is the distance from w to the nearest point of the extended medial axis.*

The LFS is a measure of how close another part of the surface is to a surface point. Note that the LFS of x can be smaller than the minimal distance to the centers of the maximal balls touching in x.

Definition 22. *For $r > 0$, a set $X \subseteq W$ is an* r-sample *if the distance from any point $w \in W$ to its closest neighbor in X is at most $r \cdot \mathrm{LFS}(w)$*

For an r-sample with $r \leq 0.1$, the algorithm labels the polar balls correctly. Furthermore, the union of the inner polar balls is homeomorphic to B_W. For the proofs in the next paragraph we assume that this sampling condition is satisfied.

Construction of an ε-admissible set We will use the power crust to obtain an ε-admissible set from an r-sample X of points on a surface W. Let $\mathcal{P}$ be the set of inner polar balls and $M = \cup\mathcal{P}$.

Condition (i) (see page 35) is proven directly in [6, Lemma 25], that states that there exists a homeomorphism between ∂M and W. For condition (ii), we prove the following lemma.

Lemma 23. *All inner polar balls $\mathcal{P}$ are maximal balls in M.*

Proof. The maximal balls in $\mathcal{P}$ are polar balls of X. These polar balls are centered at Voronoi vertices of X. Therefore, each $\hat{p} \in \mathcal{P}$ touches four points in X. Since all points in X lie on ∂M [6, Observation 3], all input balls are maximal. □

Although all balls in $\mathcal{P}$ are maximal in ∂M, there can be weighted points that touch ∂M only in intersection curves or intersection points of other weighted points.

For any $s < 1$, smooth patches grow over the intersection curves and these weighted points are not maximal anymore. We do not consider these weighted points if we compute the minimal shrink factor for which all balls are maximal.

We prove the condition (iii) using the tubular neighborhood of the surface W, cf. [6].

Definition 24. *The tubular neighborhood around the surface W is the set of points within distance $r \cdot \mathrm{LFS}(w)$ of a point $w \in W$*

We can now prove criterion (iii).

Lemma 25. *The Hausdorff distance $d(B_W, M)$ is $O(r)$.*

Proof. The bodies M and B_W only differ in the tubular neighborhood T of W, since all poles are labeled correctly ([6, Lemma 37]) and their boundaries lie in the tubular neighborhood ([6, Theorem 21] and the definition of T).

Consider a point $m \in M \cap T$, the distance of m to $B_W \setminus T$ is $O(r)\,\mathrm{LFS}(w)$ for some $w \in W$. We bound $\mathrm{LFS}(w)$ by $\max_{w' \in W} \mathrm{LFS}(w')$, which is a constant since W does not change. Thus the distance of m to B_W is $O(r)$. A similar reasoning holds for the distance between a point $w \in B_W \cap T$ and M. □

To conclude, the power crust constructs an ε-admissible set of maximal balls from an r-sample on a surface, where $\varepsilon = O(r)$. For small ε this means that there is a constant c such that $\varepsilon \leq c \cdot r$. Therefore, the power crust can construct an ε-admissible set from an ε/c-sample. The constant c can be rather large. This is due to the fact that the Hausdorff distance is a global measure between the two surfaces, whereas the local feature size is defined locally on the surface.

3.7 Conclusion and future work

We presented an algorithm to effectively compute a C^1-approximation S of a C^2-surface W represented by a set of approximate W-maximal balls. The approximation S is a skin surface, which is homeomorphic to the boundary of the union of the approximate W-maximal balls. Furthermore, the Hausdorff distance between the regions enclosed by W and S converges to zero as we increase the density of the sample of maximal balls.

A disadvantage of our method is that the surface is usually bumpy, i.e., the error of the tangent vector (the C^1-error) is not bounded, since the interpolating patches between balls are always concave. Another drawback of our algorithm is that it determines the shrink factor globally: if a high shrink factor is needed to satisfy some condition in a small part of the skin surface, this influences the approximation of the whole surface. In the next chapter we show that it is possible to use different values of the shrink factor in different parts of the surface, according to some criterion.

For s close to 1, the skin surface and the boundary of the union of the input balls is almost the same. This would imply that our approach hardly improves on the union of the balls. We assume that the shrink factor will be significantly smaller

than 1. Let s_W be the shrink factor obtained by our algorithm for the medial axis transform of W, we conjecture:

Conjecture 26. *For a C^2-surface W there is an $s_W < 1$ such that the shrink factor s_S corresponding to a sample S of the Medial Axis Transform of W converges to s_W if the sampling density goes to 1.*

We think that the value s_W depends upon the maximal ratio between the radius of the maximal ball touching a point $x \in W$ and the local feature size in x. The larger this ratio, the higher s_W.

Chapter 4

Envelope surfaces

4.1 Introduction

In this chapter we construct a closed C^1 manifold in $\mathbb{R}^d$ that wraps tightly around the union of a set of balls. Such a manifold is useful since the union of a finite set of balls can approximate any closed smooth manifold arbitrarily close. However, the boundary of the union of balls is not tangent continuous along the intersection curve of the boundary of two balls. The main idea is to construct a radius function defined on the convex hull of the centers that interpolates the radius of the input balls on their centers. The approximating manifold is the envelope (the boundary of the union) of the infinite set of balls defined by the radius function. Under certain conditions on the radius function the approximating manifold is C^1.

The theory of envelope manifolds holds in general for $\mathbb{R}^d$, but their main application is in $\mathbb{R}^2$ and $\mathbb{R}^3$. Therefore we write envelope curve or envelope surface even if the results apply manifolds in $\mathbb{R}^d$.

In the previous chapter we developed a method for approximating a surface in a similar way using skin surfaces. The approximating surface obtained by the approximation algorithm presented in Chapter 3 has nice properties:

- the surface is tangent continuous,
- it is homeomorphic to the union of the input balls,
- every ball is tangent to the boundary of the union of the input balls also touches the approximating surface and
- the Hausdorff distance between the boundary of the union of input balls and the approximating surface is bounded by an arbitrarily small, positive constant.

The main drawback of this method is that the patches interpolating between balls are concave and that the approximation is controlled by a global parameter.

We generalize the theoretical framework of skin surfaces by introducing the much larger class of *envelope surfaces*, more precisely envelopes of balls. Our scheme allows

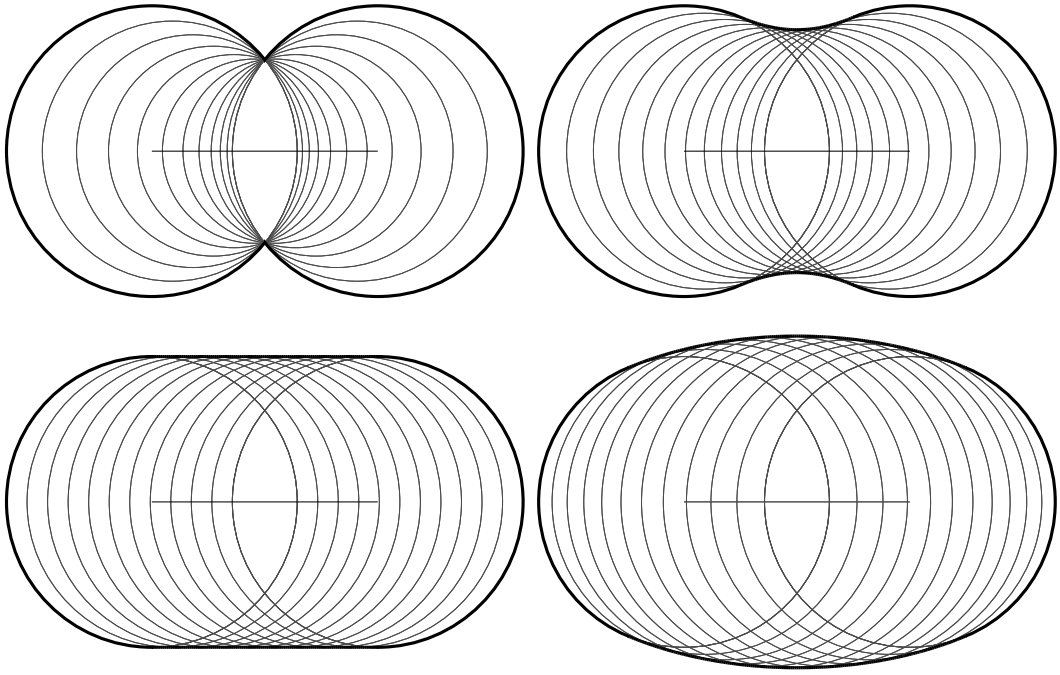

Figure 4.1: An example of an envelope surface. The weight function, defined on a line segment, is a quadratic univariate function with leading coefficient 1, 0.5, 0 and -0.5, respectively (from left to right and top to bottom).

for a local control of the surface. Therefore it is possible to interpolate between certain balls with convex patches while the interpolating patch between other balls is concave.

General idea A ball $B(p, r)$ is defined by its center p and radius r. A set of balls corresponds to a point set with a radius assigned to each point. The idea is to construct a continuous radius function ρ defined on a convex and compact domain D containing the point set, and interpolating the radii of a given finite set of balls with centers in D. We define the approximating surface S as the boundary of the infinite set of balls $\{B(p, \rho(p)) \mid p \in D\}$.

For the analysis of the surface it is more convenient to construct a squared radius function, also called weight function W. Hence, $\rho(\cdot) = \sqrt{W(\cdot)}$. If $W(p)$ is negative, then $\rho(p)$ is imaginary and the ball centered at p is omitted from the infinite set of balls.

The approximating surface contains the input balls if the radius function interpolates the radii on points of the point set. Moreover, ρ is bounded since it is continuous and defined on a compact domain. The approximating surface is therefore closed and compact. However, S is not necessarily C^1. Under certain conditions on the weight function W, the surface S is C^1, as is shown in Section 4.3.

In Section 4.4, we let the weight function be a piecewise quadratic function on a triangulation of the point set, and let D be the convex hull of the input points. In this setting, the approximating surface S is decomposable into pieces of quadrics.

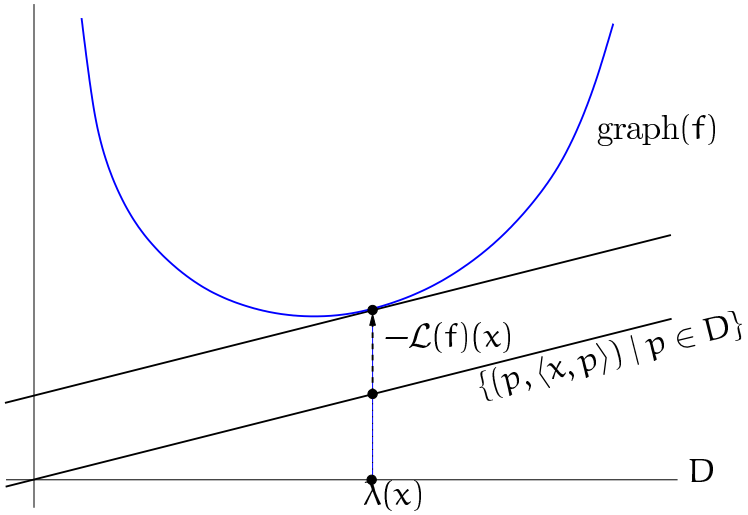

Figure 4.2: The function $\mathcal{L}(f)$ is the Legendre-Fenchel transform (conjugate) of the convex function f.

We construct a polyhedral complex that partitions $\mathbb{R}^d$ into a finite set of polyhedral cells such that the intersection of the surface S and a cell is a piece of a quadric. We conclude the chapter by showing how skin surfaces fall in this framework and how envelope surfaces can be used in surface design.

An example in 2D of an envelope curve is presented in Figure 4.1, where the set of centers D is the line segment connecting the centers of two input disks. The weight function is quadratic. Different choices of the quadratic coefficient c yield different interpolating curves as shown in the pictures: the union of the circles $c = 1$, a concave interpolating patch for $0 < c < 1$, a linear interpolation for $c = 0$ and a convex interpolating patch for $c < 0$. Examples of envelope surfaces that are not C^1 are given in Figures 4.1(a), 4.3 and 4.4.

4.2 The Legendre-Fenchel transform

We use the Legendre-Fenchel transform to show under which conditions the envelope surface is C^1. We first derive some of its properties.

Let $f : D \to \mathbb{R}$ be a continuous function defined on a compact subset D of $\mathbb{R}^d$.

Definition 27. *The* Legendre-Fenchel transform (or: conjugate) *of* f *is the function* $\mathcal{L}(f) : \mathbb{R}^d \to \mathbb{R}$, *defined by*

$$\mathcal{L}(f)(x) = \max_{p \in D}(\langle x, p\rangle - f(p)). \tag{4.1}$$

See Figure 4.2. Note that the maximum exists, since f is continuous and D is compact.

Examples. 1. Let $D = [-1, 1]$, and let $f : D \to \mathbb{R}$ be defined by $f(p) = p^2$. Then

$$\mathcal{L}(f)(x) = \begin{cases} -x - 1, & \text{if } x \leq -2, \\ \frac{1}{4}x^2, & \text{if } -2 \leq x \leq 2, \\ x - 1, & \text{if } x \geq 2. \end{cases}$$

2. Let $D = [a, b]$, with $a < 0 < b$, and let $f : D \to \mathbb{R}$ be defined by $f(p) = \mid p \mid$. Then

$$\mathcal{L}(f)(x) = \begin{cases} a(x+1), & \text{if } x \leq -1, \\ 0, & \text{if } -1 \leq x \leq 1, \\ b(x-1), & \text{if } x \geq 1. \end{cases}$$

Let $f : D \to \mathbb{R}$ be a strictly convex function defined on a compact subset D of $\mathbb{R}^d$. For $x \in \mathbb{R}^d$ there is a unique point, denoted by $\lambda(x)$, at which the maximum in (4.1) is attained. According to Lemma 64, the function $\lambda : \mathbb{R}^d \to D$ is continuous.

Lemma 28. *For $x \in \mathbb{R}^d$ there is a unique point, denoted by $\lambda(x)$, at which the maximum in (4.1) is attained. The function $\lambda : \mathbb{R}^d \to D$ is continuous, and the set $\lambda^{-1}(p)$ is convex, for $p \in D$.*

Proof. Uniqueness of $\lambda(x)$ follows from the strict convexity of f (See also Figure 4.2). According to Lemma 64, the function $\lambda : \mathbb{R}^d \to D$ is continuous.

To show that $\lambda^{-1}(p)$ is convex, let $x_0, x_1 \in \mathbb{R}^d$ such that $\lambda(x_0) = \lambda(x_1) = p_0$. From the strict convexity of f it follows that

$$\langle x, p \rangle - f(p) < \mathcal{L}(f)(x), \text{ for } p \neq \lambda(x).$$

Let $x' = (1-\gamma)x_0 + \gamma x_1$, with $0 \leq \gamma \leq 1$, be a point on the line segment x_0x_1. Then, for $p \neq p_0$:

$$\begin{aligned} \langle x', p \rangle - f(p) &= (1-\gamma)\left(\langle x_0, p \rangle - f(p)\right) + \gamma\left(\langle x_1, p \rangle - f(p)\right) \\ &< (1-\gamma)\,\mathcal{L}(f)(x_0) + \gamma\,\mathcal{L}(f)(x_1) \\ &= (1-\gamma)\left(\langle x_0, p_0 \rangle - f(p_0)\right) + \gamma\left(\langle x_1, p_0 \rangle - f(p_0)\right) \\ &= \langle x', x_0 \rangle - f(p_0). \end{aligned}$$

Therefore, $\lambda(x') = p_0$, so $\lambda^{-1}(p_0)$ is convex. □

Note that the set $\lambda^{-1}(p)$ is connected since it is convex.

Proposition 29. *If $f : D \to \mathbb{R}$ is a strictly convex continuous function defined on a compact subset D of $\mathbb{R}^d$, then the Legendre-Fenchel transform $\mathcal{L}(f)$ is a convex C^1-function. Its derivative at $x \in \mathbb{R}^d$ is the linear function*

$$\mathcal{L}(f)'(x) = \langle \lambda(x), \cdot \rangle. \tag{4.2}$$

A point p_0 *is equal to* $\lambda(x_0)$ *if and only if the directional derivatives (cf. Section A.2) are non-positive:*

$$\frac{\partial \Phi}{\partial p}(x_0, p_0)(v) \leq 0, \text{ for all } v \in \mathbb{R}^d,$$

where the function $\Phi : \mathbb{R}^d \times D \to \mathbb{R}$ *is defined by* $\Phi(x, p) = \langle x, p \rangle - f(p)$.

If f *is a* C^k*-function, with* $k \geq 2$*, on a neighborhood of* $p_0 \in D$*, then:*

- $\frac{\partial \Phi}{\partial p}(x, \lambda(x))(v) = 0$ *for all* $v \in \mathbb{R}^d$ *and*
- *there is a unique point* $x_0 \in \mathbb{R}^d$ *with* $p_0 = \lambda(x_0)$ *and*
- $\mathcal{L}(f)$ *is a* C^k *function on a neighborhood of* x_0.

Furthermore, $p = \lambda(x)$ *is the unique solution of the equation*

$$x = \nabla f(p),$$

for x *near* x_0 *and* p *near* p_0.

Proof. For $p \in D$ the function $\Phi(\cdot, p)$ is affine, hence convex. Since $\mathcal{L}(f)(x) = \max_{p \in D} \Phi(x, p)$, Lemma 68 implies that the Legendre-Fenchel transform $\mathcal{L}(f)$ is convex.

We first prove that the Legendre-Fenchel transform $\mathcal{L}(f)$ is Gateaux-differentiable on $\mathbb{R}^d$. To this end, let $x, v \in \mathbb{R}^d$, and let $h > 0$. It follows from the definition of Φ that

$$\frac{\mathcal{L}(f)(x + hv) - \mathcal{L}(f)(x)}{h} \geq \frac{\Phi(x + hv, \lambda(x)) - \Phi(x, \lambda(x))}{h} = \langle v, \lambda(x) \rangle.$$

Since $\mathcal{L}(f)$ is convex, its directional derivative $\mathcal{L}(f)'(x; v)$, cf. Section A.2, exists according to Lemma 69. The last identity yields, after taking the limit $h \downarrow 0$:

$$\mathcal{L}(f)'(x; v) \geq \langle v, \lambda(x) \rangle. \tag{4.3}$$

To prove the reverse inequality observe that convexity of $\mathcal{L}(f)$ implies

$$\mathcal{L}(f)'(x; v) \leq \frac{\mathcal{L}(f)(x + hv) - \mathcal{L}(f)(x)}{h}.$$

Hence

$$\begin{aligned} \mathcal{L}(f)'(x; v) &\leq \frac{\Phi(x + hv, \lambda(x + hv)) - \Phi(x, \lambda(x))}{h} \\ &\leq \frac{\Phi(x + hv, \lambda(x + hv)) - \Phi(x, \lambda(x + hv))}{h} \\ &= \langle v, \lambda(x + hv) \rangle \end{aligned}$$

Using the continuity of λ (cf. Lemma 64) we obtain, after taking the limit $h \downarrow 0$:

$$\mathcal{L}(f)'(x;v) \leq \langle v, \lambda(x) \rangle. \tag{4.4}$$

From (4.3) and (4.4) we conclude that $\mathcal{L}(f)$ has a directional derivative at $x \in \mathbb{R}^d$ in all directions $v \in \mathbb{R}^d$, and that

$$\mathcal{L}(f)'(x;v) = \langle v, \lambda(x) \rangle.$$

Therefore all partial derivatives of $\mathcal{L}(f)$ at x exist, and are continuous. Hence $\mathcal{L}(f)$ is even Fréchet differentiable with derivative $\mathcal{L}(f)'(x) = \langle \lambda(x), \cdot \rangle$. Since this derivative is continuous, $\mathcal{L}(f)$ is a C^1-function.

All directional derivatives of $\Phi(x, \cdot)$ exist by Lemma 69. Since the maximum of $\Phi(x, \cdot)$ is attained at $\lambda(x)$ we have

$$\frac{\partial \Phi}{\partial p}(x, \lambda(x))(v) \leq 0, \text{ for all } v \in \mathbb{R}^d$$

Conversely, if $\frac{\partial \Phi}{\partial p}(x, p_0)(v) \leq 0$ for all $v \in \mathbb{R}^d$, then $p_0 = \lambda(x)$ since $\Phi(x, \cdot)$ is strictly convex.

If f is C^k at $\lambda(x)$, then the point $\lambda(x)$ is a critical point of the function $\Phi(x, \cdot)$, so $\nabla_p \Phi(x, \lambda(x)) = 0$. In other words, $p = \lambda(x)$ is the solution of the equation

$$x - \nabla f(p) = 0.$$

So let $x_0 = \nabla f(p_0)$, then $\lambda(x_0) = p_0$. Since $k \geq 2$ and f is strictly convex, the Hessian of f exists and is positive definite on a neighborhood V of p_0. In other words, the C^{k-1}-function $\nabla f : \mathbb{R}^d \to \mathbb{R}^d$ is locally invertible at p_0. Its inverse is the function λ, restricted to the neighborhood V of p_0. The singleton set $\{x_0\}$ is a connected component of $\lambda^{-1}(p_0)$. Since the latter set is connected according to Lemma 28, we conclude that

$$\lambda^{-1}(p_0) = \{x_0\}.$$

In other words, x_0 is the unique point with $p_0 = \lambda(x_0)$. By the Inverse Function Theorem, λ is also a C^{k-1}-function. In view of (4.2) the derivative of $\mathcal{L}(f)$ is a C^{k-1}-function on V, so $\mathcal{L}(f)$ is a C^k-function on V. □

Remark. As can be seen from the second example preceding Proposition 29, the result does not necessarily hold if the function f is not strictly convex. In that case f is convex and λ is discontinuous.

4.3 Envelope surfaces

Let D be a compact convex subset of $\mathbb{R}^d$, and let $W : D \to \mathbb{R}$ be a continuous function. Consider the family of spheres $\{C_p \mid p \in D\}$, where C_p is the sphere with center p and weight $W(p)$. A point $x \in \mathbb{R}^d$ lies on the *envelope* S of this family of

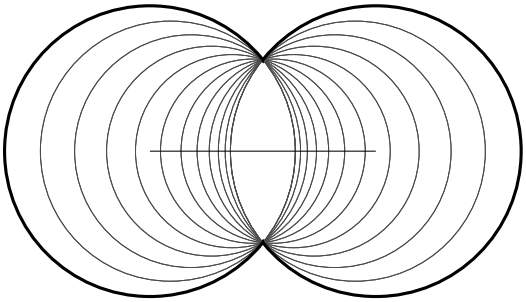

Figure 4.3: An example where the weight function is chosen so that the envelope surface is the boundary of the union of the input circles.

spheres if it lies on at least one sphere, and on or outside all other spheres. In other words, x lies on the envelope if $H(x) = 0$, where $H : \mathbb{R}^d \to \mathbb{R}$ is defined by

$$H(x) = \min_{p \in D} \left(\| x - p \|^2 - W(p) \right). \tag{4.5}$$

Since D is compact and W is continuous, the minimum is attained in at least one point of D. The main result of this section states that, under certain conditions on the weight function W, the zero set of H is a C^1-submanifold of $\mathbb{R}^d$.

Proposition 30. *Assume that $W_1 : D \to \mathbb{R}$, defined by $W_1(p) = \| p \|^2 - W(p)$, is a strictly convex function. Then*

1. *For $x \in \mathbb{R}^d$ the minimum in (4.5) is attained in a unique point $\lambda(x)$, and the function $\lambda : \mathbb{R}^d \to D$ is continuous.*

2. *H is a C^1-function, with derivative*

$$H'(x) = 2\langle x - \lambda(x), \cdot \rangle$$

   In particular, x is a critical point of H if and only if $x \in D$ and $x = \lambda(x)$.

3. *The zero set of H is a C^1-submanifold of $\mathbb{R}^n$ if it does not contain any zeros of W.*

4. *If W_1 is a C^k-function in a neighborhood of p_0, with $k \geq 1$, then there is a unique point x_0 with $p_0 = \lambda(x_0)$, and H is a C^k-function on a neighborhood of x_0.*

Proof. 1. Rewrite H as follows

$$\begin{aligned} H(x) &= \| x \|^2 - \max_{p \in D} \left(2\langle p, x \rangle - \| p \|^2 + W(p) \right) \\ &= \| x \|^2 - 2\mathcal{L}(1/2\, W_1)(x). \end{aligned} \tag{4.6}$$

Here $\mathcal{L}(1/2\, W_1) : \mathbb{R}^d \to \mathbb{R}$ is the Legendre-Fenchel transform (or: conjugate) of $1/2\, W_1$. We refer to Section 4.2 for the definition, and some key properties relevant

for our context. Since W_1 is strictly convex, for $x \in \mathbb{R}^n$ there is a unique point $\lambda(x)$ in D in which the maximum in (4.6) is attained. Obviously, the minimum in (4.5) is also attained at $\lambda(x)$. According to Lemma 64, the function $\lambda : \mathbb{R}^d \to D$ is continuous.
2. According to Proposition 29 the $\mathcal{L}(1/2\,W_1)$ is C^1, since D is compact and W_1 is strictly convex. Therefore H is a C^1-function, the derivative of which satisfies

$$H'(x) = 2\langle x, \cdot\rangle - 2\mathcal{L}(1/2\,W_1)'(x) = 2\langle x - \lambda(x), \cdot\rangle.$$

The expression for the derivative of $\mathcal{L}(1/2W_1)$ is given in Proposition 29.
3. From the second part we conclude that H is C^1 and we assume that $k \geq 2$. By the Implicit Function Theorem, it is sufficient to prove that H is regular at every point of its zero set. Let x be a point of $H^{-1}(0)$. Since $\lambda(x) \in D$, every point x in the complement of D is different from $\lambda(x)$. According to the second part of the proposition, H is regular at such points. Next consider $x \in D$. Since $W(x) \neq 0$, it follows from

$$\|x - \lambda(x)\|^2 - W(\lambda(x)) = H(x) = 0$$

that $x \neq \lambda(x)$. Therefore, H is regular at x.
4. This follows directly from Proposition 29. □

Remark. If W_1 is not strictly convex, then the zero set of H need not be a C^1-surface. An example is given in Figure 4.3. The weight function is defined on a line segment and is given by $W(p) = \|p\|^2 + c$, with $c > 0$. The function W_1 is constant and therefore convex, but not strictly convex. In this case the zero set of H is not C^1.

Remark. If $\lambda(x)$ is an interior point of D, and W is differentiable at $\lambda(x)$, then $x \in H^{-1}(0)$ if there is a point p (namely, $p = \lambda(x)$), such that

$$\Phi(x, p) = 0 \quad \text{and} \quad \nabla_p \Phi(x, p) = 0.$$

Here $\Phi : \mathbb{R}^d \times D \to \mathbb{R}$ is defined by $\Phi(x, p) = \|x - p\|^2 - W(p)$. Note that $\Phi(x, p) = \|x\|^2 - 2\,f(x, p)$, with $f(x, p) = \langle x, p\rangle - 1/2\,W_1(p)$. Let $x \in H^{-1}(0)$ and $x \neq \lambda(x)$, then $\Phi(x, p) = 0$ and $\Phi(x, \cdot)$ has a maximum at $\lambda(x)$, even if W is not differentiable at $\lambda(x)$ or if $\lambda(x)$ lies on the boundary of D. Denote by Σ the set:

$$\Sigma = \{x \in \mathbb{R}^d \mid \Phi(x, p) = 0 \text{ and } \Phi(x, \cdot) \text{ has a maximum at } p, \text{ for } p \in D\}. \qquad (4.7)$$

If $x \neq \lambda(x)$ for all $x \in D$, then the envelope surface $H^{-1}(0)$ is contained in Σ. The envelope surface and Σ are equal if W_1 is strictly convex. If W_1 is not strictly convex then $\Phi(x, \cdot)$ may have more than one critical point and $H^{-1}(0)$ may be a proper subset of Σ, as is shown in Figure 4.4.

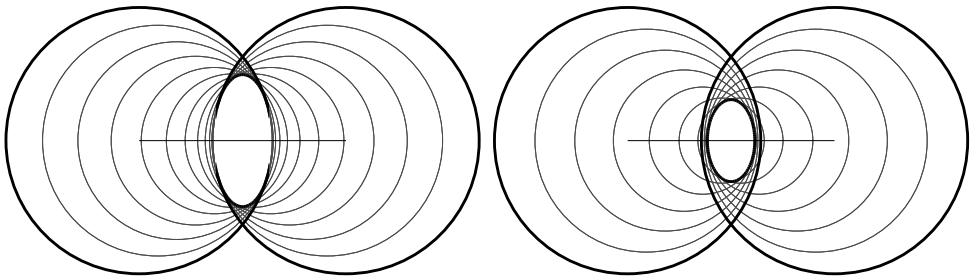

Figure 4.4: In these two cases the envelope surface is a proper subset of Σ defined by Equation (4.7). The weight function $W : [-1, 1] \to \mathbb{R}$ is defined by $W(p) = 1.25p^2 + 0.75$ (left) and by $W(p) = 1.5p^2 + 0.5$ (right).

Piecewise smooth envelopes Piecewise smooth envelope surfaces can be constructed from piecewise smooth weight functions. So let D be a compact convex polyhedron in $\mathbb{R}^d$, endowed with a triangulation T. The function $W : D \to \mathbb{R}$ is piecewise C^k with respect to T if

1. W is continuous on D, and
2. the restriction W_σ of W to a d-dimensional simplex σ of T can be extended to a C^k-function on a neighborhood of the closure of σ.

The envelope surface associated with the family of spheres with center p and weight $W(p)$, with p ranging over D, is a C^1 surface if the function $W_1 : D \to \mathbb{R}$, defined by

$$W_1(p) = \|p\|^2 - W(p), \tag{4.8}$$

is strictly convex. Each of the functions W_σ should therefore satisfy the following condition.

Strict Convexity Condition. The function $W_{\sigma,1}$, defined by $W_{\sigma,1}(p) = \|p\|^2 - W_\sigma(p)$ on the domain of W_σ, is strictly convex.

Obviously, this condition does not guarantee strict convexity of W_1 at lower dimensional simplices of T. Therefore we require that the family $\{W_\sigma\}$ satisfies the following condition.

Monotonous Transition Condition. If the d-dimensional simplices σ and σ' share a facet τ with normal ν directed from σ to σ', then for $p_0 \in \tau$:

$$W'_{\sigma'}(p_0; \nu) \leq W'_\sigma(p_0; \nu).$$

Lemma 31. *Let $W : D \to \mathbb{R}$ be a piecewise C^1-function with respect to the triangulation T, satisfying the Strict Convexity Condition and the Monotonous Transition Condition. Then the function $W_1 : D \to \mathbb{R}$, defined by (4.8), is strictly convex.*

Proof. We shall prove that W_1 is strictly convex along line segments $l \cap D$, where l is a line in $\mathbb{R}^d$. To this end, it is sufficient to prove that W_1 is convex along such

lines. Indeed, the triangulation T partitions $l \cap D$ in a finite alternating sequence of points and open intervals, formed by the intersection of l with the simplices of T. Let $l \cap \tau$ be such an open interval, and suppose the simplex τ lies in the closure of the d-dimensional simplex σ. Then W_l coincides with W_σ on this interval, so $W_l|_{l\cap\tau}$ is strictly convex. Therefore Lemma 66 in Appendix A.3 allows us to conclude that W_l is strictly convex once we know that W_l is convex along l (or, more precisely, that $W_l|_{l\cap D}$ is convex).

We first prove that W_l is convex on any line disjoint from the $(d-2)$-skeleton of T, and conclude, via a limit process, that W is convex on any line.

So consider a line l with direction v that only intersects simplices of dimension d and $d-1$. Since W satisfies the Strict Convexity Condition, it is sufficient to prove that

$$W'_{\sigma',l}(p_0; v) \leq W'_{\sigma,l}(p_0; v),$$

for a point p_0 on a $(d-1)$-simplex τ shared by two d simplices σ and σ', such that v is directed into σ. Note that v is not parallel to τ since l is disjoint from the $(d-2)$-skeleton of T. Write $v = v_{\|} + v_{\perp}$, where $v_{\|}$ is parallel to τ and $v_{\perp}$ is perpendicular to τ. Since $W_{\sigma,l}$ and $W_{\sigma',l}$ are the restrictions of C^1-functions defined on a neighborhood of p_0 in $\mathbb{R}^d$, the derivatives of these functions at p_0 exist and are linear, so:

$$W'_{\sigma,l}(p_0; v) = W'_{\sigma,l}(p_0; v_{\|}) + W'_{\sigma,l}(p_0; v_{\perp}).$$

and

$$W'_{\sigma',l}(p_0; v) = W'_{\sigma',l}(p_0; v_{\|}) + W'_{\sigma',l}(p_0; v_{\perp}).$$

Since the restrictions of the functions $W_{\sigma,l}$ and $W_{\sigma',l}$ to the facet τ are equal, the derivatives of these functions at p_0 in directions parallel to v are equal. In particular

$$W'_{\sigma',l}(p_0; v_{\|}) = W'_{\sigma,l}(p_0; v_{\|}).$$

According to the Monotonous Transition Condition

$$W'_{\sigma',l}(p_0; v_{\perp}) \leq W'_{\sigma,l}(p_0; v_{\perp}).$$

These identities and inequalities imply that

$$W'_{\sigma',l}(p_0; v) \leq W'_{\sigma,l}(p_0; v).$$

Hence W_l is convex along l.

To prove that W_l is convex along any line l, let l_n, $n = 0, 1, \ldots$, be a sequence of lines converging to l, such that l_n is disjoint from the $(d-2)$-skeleton of T. According to Appendix A, Lemma 65, the restriction of W_l to l is a convex function, since it is the limit of the sequence of convex functions $W_l|_{l_n}$. □

If the function W is piecewise C^k with respect to the triangulation T of D and W_l is strictly convex, then the function H is piecewise C^k on the pull-back $\lambda^*(T)$ of T under $\lambda : \mathbb{R}^d \to D$. The pull-back $\lambda^*(T)$ is the subdivision of $\mathbb{R}^d$ into the regions

$\lambda^{-1}(\tau)$, where τ ranges over all simplices of T. In general this subdivision can have a weird structure, even though the regions $\lambda^{-1}(\tau)$ are *connected*. However, if W is *piecewise quadratic* with respect to T, the cells of the pull-back are polyhedra. This context is studied in detail in the next section.

4.4 Piecewise quadratic weight functions

Suppose that we have a finite set of balls $\mathcal{P}$ and a triangulation T of their centers. In accordance with the ideas presented in the previous section, we construct an envelope surface by defining a piecewise quadratic weight function on T. This weight function is differentiable except possibly at lower dimensional cells, where it is merely assumed to be continuous.

If the weight function is interpolating, then the envelope surface contains the input balls. In this section we show that the envelope surface is piecewise quadratic and that there is a polyhedral complex partitioning the envelope surface into the quadratic pieces.

Quadratic functions A quadratic function q on $\mathbb{R}^d$ is of the form

$$q(x) = x^T Q x + a^T x + b,$$

for a symmetric $d \times d$ matrix Q, a d-vector a and a constant b. The matrix Q is called the defining matrix of q. Since Q is symmetric, it has real eigenvalues [71, Ch. 7]. If the eigenvectors of Q span the coordinate axes, then Q is a diagonal matrix with the eigenvalues on the diagonal.

Quadratic weight function In this section, the weight function is piecewise quadratic. Each quadratic function is defined on a simplex. We analyze the envelope surface corresponding to a single quadratic weight function. In order to avoid boundary conditions, we extend the domain D of the weight function to an n-dimensional affine subspace of $\mathbb{R}^d$ containing the domain. Hence we extend the line segment of centers in Figure 4.1 to the line containing the segment and avoid spherical patches where the outer circles touch the envelope surface.

These weight functions are of special interest since the corresponding envelope surfaces are quadrics. Note that in this setting, D is not compact as assumed in the text above. Therefore, the minimum in Equation (4.5) might not be attained. Write

$$\begin{aligned} H(x) &= \inf_{p \in D} \left(\|x - p\|^2 - W(p) \right) \\ &= \|x\|^2 - \inf_{p \in D} \left(2\langle p, x \rangle - W_1(p) \right), \end{aligned}$$

where $W_1 : D \to \mathbb{R}$ is defined by $W_1(p) = \|p\|^2 - W(p)$. If W_1 is strictly convex, then $2\langle p, x \rangle - W_1(p)$ has a positive definite quadratic form and the infimum is attained. Hence,

$$H(x) = \min_{p \in D} \left(\|x - p\|^2 - W(p) \right).$$

Lemma 32. *If W is a quadratic function, defined on an affine subspace D of $\mathbb{R}^d$ and W_1 defined by $W_1(p) = \|p\|^2 - W(p)$, is a strictly convex function, then*

1. *The map λ is linear,*
2. *The set $\lambda^{-1}(p)$ is orthogonal and complementary to D.*
3. *The function H is a quadratic function that is rotationally symmetric in D,*
4. *The eigenvectors of W and H restricted to D coincide and an eigenvalue d_i of W corresponds to an eigenvalue $d_i/(d_i-1)$ of H. The eigenvalues of H in the direction orthogonal to D are 1. The critical point of W and H are equal and the critical values in the apex are equal with opposite signs.*

Proof. 1. Let $W : D \to \mathbb{R}$ be given by $W(p) = p^T Q p + a^T p + c$, where Q is a symmetric positive definite $n \times n$-matrix with eigenvalues less than one, a is a vector in $\mathbb{R}^n$ and c is a real constant. Let $\alpha : \mathbb{R}^d \to D$ be the affine map projecting a point orthogonally on D

Then $\Phi(x,p) = \|x - \alpha(x)\|^2 + \tilde{\Phi}(\alpha(x),p)$, where $\tilde{\Phi} : D \times D \to \mathbb{R}$ is the restriction of Φ to $D \times D$. Hence $\nabla_p \Phi(x,p) = \nabla_p \tilde{\Phi}(\alpha(x),p) = -2\alpha(x) - 2(Q-I)p - a$. So $\lambda : \mathbb{R}^d \to \mathbb{R}^n$ is given by:

$$\lambda(x) = -(Q-I)^{-1}(\alpha(x) + \frac{1}{2}a). \tag{4.9}$$

So, $\lambda(x)$ depends linearly on x.

2. Furthermore, $\lambda^{-1}(p)$ is the pre-image of the point $-(Q-I)p - 1/2a \in D$ under the orthogonal projection from $\mathbb{R}^n$ onto $\mathbb{R}^d$. Therefore, $\lambda^{-1}(p)$ is orthogonal and complementary to D.

3. Since

$$H(x) = \|x - \lambda(x)\|^2 - W(\lambda(x)),$$

it follows that H depends quadratically on x. Finally,

$$H(x) = \|x - \alpha(x)\|^2 + \|\alpha(x) - \lambda(x)\|^2 - W(\lambda(x)). \tag{4.10}$$

If x_1 and x_2 are in an affine subspace orthogonal to D and at the same distance from D, then $\alpha(x_1) = \alpha(x_2)$, $\|x_1 - \alpha(x_1)\| = \|x_2 - \alpha(x_2)\|$, and, by the second part of the lemma $\lambda(x_1) = \lambda(x_2)$. Therefore $H(x_1) = H(x_2)$, so H is rotationally symmetric with respect to D.

4. It follows from (4.10) that H is a quadratic function with matrix of the form

$$R = \begin{pmatrix} \tilde{R} & 0 \\ 0 & I \end{pmatrix},$$

with respect to the orthogonal splitting of $\mathbb{R}^d$ into spaces parallel and orthogonal to D. To determine the eigenvalues of $\tilde{R}$, observe that according to (4.9): $\alpha(x) = (I-Q)\lambda(x) - 1/2a$, so:

$$\begin{aligned} & \| \alpha(x) - \lambda(x) \|^2 - W(\lambda(x)) \\ = \; & \| -Q\lambda(x) - 1/2a \|^2 - \lambda(x)^T Q \lambda(x) - a^T \lambda(x) - c \\ = \; & \lambda(x)^T (Q^2 - Q)\lambda(x) + a^T (Q - I)\lambda(x) + 1/4 a^T a - c \\ = \; & \alpha(x)^T (Q-I)^{-1} Q \alpha(x) + a^T (Q-I)^{-1} \alpha(x) - 1/4 \| a \|^2 - c. \end{aligned}$$

Therefore $\tilde{R} = Q(Q-I)^{-1}$, so the eigenvalues of $\tilde{R}$ are $d_i/(d_i-1)$, with $1 \leq i \leq n$.

The critical point of W is the point $p_0 \in D$ such that $\nabla W(p_0) = 0$, hence

$$p_0 = -\frac{1}{2} Q^{-1} a.$$

Since H is symmetric in D by part 3, the critical point of H lies on D. The gradient of H vanishes at p_0 if $2(Q-I)^{-1} Q p_0 + (Q-I)^{-1} a = 0$, i.e., if $p_0 = -1/2 Q^{-1} a$. The critical value of W is $c - 1/4 a^T Q^{-1} a$ and the critical value of H is $-c + 1/4 a^T Q^{-1} a$, which concludes the proof. □

Corollary 33. *The type of the quadric $H^{-1}(0)$ only depends on the eigenvalues of the quadratic weight function.*

The eigenvalues of H determine the type of quadric. In view of Lemma 32, part 4, the eigenvalues of H depend only on the eigenvalues of W. For example, let W be defined on a one-dimensional affine space with leading coefficient c. Then the envelope surface is a hyperboloid if $0 < c < 1$, a cylinder for $c = 0$ and an ellipsoid for $c < 0$.

Decomposition of the surface For $x \in S$ the ball B_x with center $\lambda(x)$ and radius-squared $W(\lambda(x))$ is called the *defining ball* of x. It is tangent to S at x. For a simplex σ of T let S_σ be the set of points of the envelope surface S such that the center of their defining balls lies on the relative interior of σ. We shall prove that $\{S_\sigma \mid \sigma \in T\}$ is a decomposition of S into quadratic patches. Even stronger, we shall construct a polyhedral decomposition $\{\mu_\sigma \mid \sigma \in T\}$ of $\mathbb{R}^d$ such that S_σ is the intersection of S with the polyhedral cell μ_σ.

More precisely, for a simplex σ of T let $\mu_\sigma = \{x \in \mathbb{R}^d \mid \lambda(x) \in \text{rel-int}(\sigma)\}$. Since λ is continuous, it maps the closure $\overline{\mu_\sigma}$ of μ_σ onto σ.

Example. Let $W(p) = 1/2 \|p\|^2$ be defined on the triangle (p_0, p_1, p_2) with $p_0 = (-1,-1)$, $p_1 = (-1,2)$ and $p_2 = (2,-1)$, viz. Figure 4.5(a). We compute the polyhedral cell $\mu_{\langle p_0 \rangle}$, which is formed by all points x where $H(x) = \Phi(x, p_0)$, since $\text{rel-int}\langle p_0 \rangle = p_0$. Hence, $\mu_{\langle p_0 \rangle} = \{x = (\xi_0, \xi_1) \mid \xi_0, \xi_1 \leq -.5\}$. Since λ changes continuously with x, it follows that $\mu_{\langle p_0, p_1 \rangle}$ is bounded by incident simplices. Let $x = (\xi_0, \xi_1)$ and $p' \in \text{rel-int}\langle p_0, p_1 \rangle$ be the point minimizing $\Phi(x, p)$ over all $p \in \text{rel-int}\langle p_0, p_1 \rangle$. Then $x \in \mu_{\langle p_0, p_1 \rangle}$ if $\Phi(x, p') = H(x)$. For $\xi_1 \leq -1/2$ we have

$\Phi(x,p') > \Phi(x,p_0)$, hence $\Phi(x,p') > H(x)$. Similarly for $\xi_1 \geq 1$ we have $\Phi(x,p') > \Phi(x,p_1)$ and for $\xi_0 > -1/2$, $\Phi(x,p') > \Phi(x,p'')$, with p'' the point minimizing $\Phi(x,p)$ over all $p \in \text{rel-int}\langle p_0,p_1,p_2\rangle$. Hence, $\mu_{\langle p_0,p_1\rangle} = (-\infty,-1/2] \times (-1/2,1)$.

The remaining cells of the polyhedral complex are determined similarly.

Remark. The set $\overline{\mu_\sigma}$ is different from the set $\{x \in \mathbb{R}^d \mid \lambda(x) \in \sigma\}$. The example in Figure 4.5(b)–(d) illustrates this. The cell $\overline{\mu_\sigma}$ with $\sigma = \langle p_0,p_1,p_3\rangle$ does not contain the cells μ_τ for $\tau = \langle p_0,p_3\rangle$ or $\tau = \langle p_0\rangle$. However, the set $\mu(\sigma) = \cup_{p\in\sigma}\mu(p)$ does contain these cells.

In the following proposition, we define the polyhedral complex and show that it decomposes the envelope surface in pieces of quadrics.

Proposition 34. *The cells $\overline{\mu_\sigma}$, with $\sigma \in T$, and their non-empty intersections form a polyhedral complex decomposing $\mathbb{R}^d$. Moreover, the cells have the property that*

1. *for $\sigma,\sigma' \in T$, with $\sigma' \leq \sigma$, the cells $\overline{\mu_\sigma}$ and $\overline{\mu_{\sigma'}}$ intersect*
2. *the intersection of $\overline{\mu_\sigma}$ with S is contained in the quadratic envelope surface obtained by extending W_σ to $\text{aff}(\sigma)$*

Proof. We use Lemma 35 to show that the cells $\overline{\mu_\sigma}$ and their non-empty intersections form a polyhedral complex. Before we can apply this lemma, we have to show that the cells $\overline{\mu_\sigma}$ are convex and partition $\mathbb{R}^d$.

We first show that the cells μ_τ with $\tau \in T$ form a partition of $\mathbb{R}^d$ and are pairwise disjoint. To this end, let $x \in \mathbb{R}^d$ and $\tau \in T$ be the simplex containing the point $\lambda(x) \in D$ in its relative interior. There may be more simplices containing $\lambda(x)$, but only one simplex contains $\lambda(x)$ in its relative interior. Hence, the cells are pairwise disjoint and their union is $\mathbb{R}^d$.

To prove that the cell μ_σ is convex, let $x_0,x_1 \in \mu_\sigma$ and $x' = (1-\gamma)x_0 + \gamma x_1$ for $\gamma \in [0,1]$. The cell μ_σ is convex if $x' \in \mu_\sigma$. Let $p' = (1-\gamma)\lambda(x_0) + \gamma\lambda(x_1)$. We shall prove that $\lambda(x') = p'$, then $x' \in \mu_\sigma$ since $p' \in \text{rel-int}(\sigma)$. To this end, let $\Phi : \mathbb{R}^d \times D \to \mathbb{R}$ be defined by $\Phi(x,p) = \langle x,p\rangle - W_1(p)$. By Proposition 29, the points p' and $\lambda(x')$ are equal iff

$$\frac{\partial\Phi}{\partial p}(x',p')(v) \leq 0, \text{ for all } v \in \mathbb{R}^d.$$

By the same proposition, $\frac{\partial\Phi}{\partial p}(x_0,\lambda(x_0))(v) \leq 0$ and $\frac{\partial\Phi}{\partial p}(x_1,\lambda(x_1))(v) \leq 0$. We expand the directional derivative to:

$$\frac{\partial\Phi}{\partial p}(x,p)(v) = \langle x,v\rangle - W_1'(p;v). \tag{4.11}$$

Since $x_0,x_1 \in \mu_\sigma$, the points $\lambda(x_0)$ and $\lambda(x_1)$ lie in the relative interior of σ. Let σ' be the simplex containing the points $\lambda(x_0) + \epsilon v$, for small positive ϵ. Then σ' also contains the points $\lambda(x_1) + \epsilon v$ for small positive ϵ. Hence, $W(\lambda(x_i) + \epsilon v) = W_{\sigma'}(\lambda(x_i) + \epsilon v)$, for $i = 0,1$. If v is parallel to $\text{aff}(\sigma)$, then the simplices σ and σ'

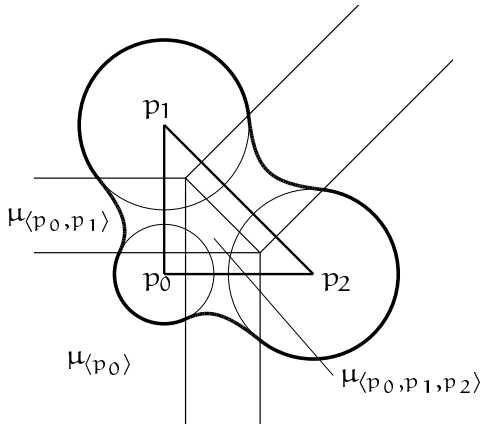

(a) Three input circles with the polyhedral complex.

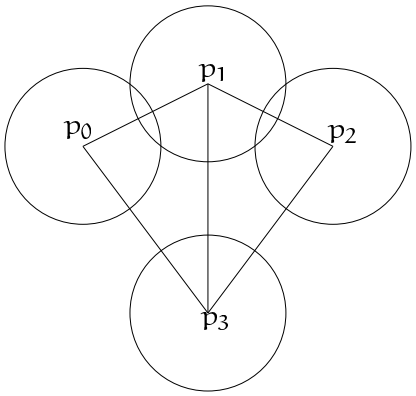

(b) The four input circles together with their triangulation.

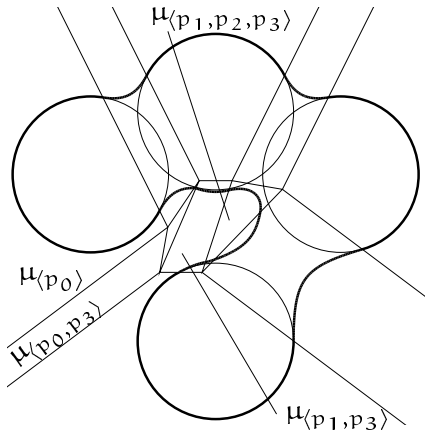

(c) The polyhedral complex decomposing the envelope surface into pieces of quadrics.

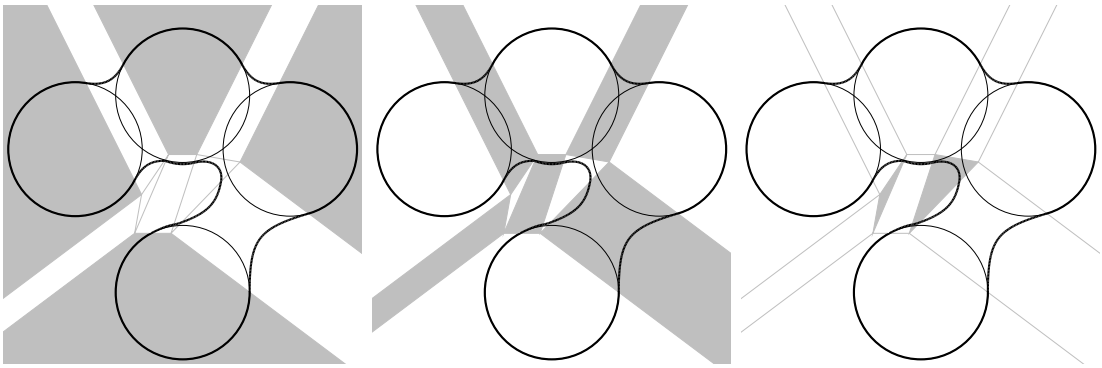

(d) The polyhedral cells μ_σ, with σ a vertex, an edge and a face.

Figure 4.5: The decomposition of the envelope surface into quadrics.

are equal, otherwise $\sigma \leq \sigma'$. Restricted to σ', W is a quadratic function. In view of (4.11), the directional derivative of Φ is linear in x and p. But then,

$$\begin{aligned}\frac{\partial \Phi}{\partial p}(x', p')(v) &= \frac{\partial \Phi}{\partial p}((1-\gamma)x_0 + \gamma x_1, (1-\gamma)\lambda(x_0) + \gamma\lambda(x_1))(v) \\ &= (1-\gamma)\frac{\partial \Phi}{\partial p}(x_0, \lambda(x_0))(v) + \gamma\frac{\partial \Phi}{\partial p}(x_1, \lambda(x_1))(v) \\ &\leq 0.\end{aligned}$$

Hence, $p' = \lambda(x')$ and $x' \in \mu_\sigma$. Therefore μ_σ is convex.

In view of Lemma 35, the cells $\overline{\mu_\sigma}$ with a non-empty interior form a polyhedral complex. A cell with empty interior is a common face of two cells with non-empty interiors by the continuity of λ.

1. Let $x \in \overline{\mu_\sigma} \setminus \mu_\sigma$ and $\tau \in T$ the simplex such that $x \in \mu_\tau$. Note that $\sigma \neq \tau$. From the continuity of λ we derive that $\lambda(x)$ is a point on the boundary of σ. Since $\lambda(x)$ lies in the relative interior of τ, we have $\tau \leq \sigma$.

2. Let x be a point in $\overline{\mu_\sigma}$. Using the continuity of λ, the point $\lambda(x)$ is contained in σ. Hence $H(x) = H_\sigma(x)$, with $H_\sigma(x) = \min_{p\in\sigma}(\|x-p\|^2 - W_\sigma(p))$. Since $W_{\sigma,1}$ is strictly convex, H has a unique minimum $\lambda(x) \in \sigma$ for all $p \in \mathrm{aff}(\sigma)$ and $H_\sigma(x) = \min_{p\in\mathrm{aff}(\sigma)}(\|x-p\|^2 - W_\sigma(p))$. Lemma 32 states that H_σ is a quadratic function. The surface S is defined as the zero set of H, hence

$$\overline{\mu_\sigma} \cap S = \overline{\mu_\sigma} \cap H^{-1}(0) = \overline{\mu_\sigma} \cap H_\sigma^{-1}(0). \qquad \square$$

The following general lemma is used in Proposition 34 to show that the cells μ_σ, with $\sigma \in T$, form a polyhedral complex.

Lemma 35. *Let $\mathcal{C}$ be a finite set of closed convex sets, with non-empty interior partitioning $\mathbb{R}^d$. The sets in $\mathcal{C}$ and their non-empty intersections form a polyhedral complex if each they only intersect in their boundaries.*

Proof. First, we show that the sets are convex polyhedra. Since the sets are convex and their interiors are disjoint, there exists a, not necessarily unique, hyperplane separating the interiors of any two sets in $\mathcal{C}$. For two sets $c, c' \in \mathcal{C}$, with $c \neq c'$, denote with $H(c, c')$ the closed half-space containing c and bounded by a hyperplane separating c and c'.

For fixed $c \in \mathcal{C}$, let x be a point in the interior of the intersection of all halfspaces $H(c, c')$, for $c' \in \mathcal{C}$ and $c' \neq c$. Since x does not lie in any set $c' \in \mathcal{C}$ for $c' \neq c$, the point x is contained in the set c. Hence,

$$c = \bigcap_{c'\in\mathcal{C}, c'\neq c} H_{c,c'},$$

which is a polyhedron by definition.

The intersection of two convex polyhedra is empty or another convex polyhedron. Since the sets in $\mathcal{C}$ only intersect in their boundary, the intersection is a face of both. $\square$

Mixed complex The polyhedral complex for envelope surfaces is a generalization of the mixed complex for skin surfaces, see Section 2.7.4. For a skin surface with shrink factor s, the graph of the weight function on a simplex σ is a paraboloid with leading term equal to s. The apex of the paraboloid lies on the intersection point of the affine hull of the Delaunay simplex δ_σ and the Voronoi cell ν_σ, which is the focus $f(\sigma)$ of σ, cf. Section 2.5.1.

A calculation in the spirit of Lemma 32 shows that $\{f(\sigma)+(1-s)x | x \in \delta_\sigma\}$ is the image of $\overline{\mu_\sigma}$ onto $\mathrm{aff}(\delta_\sigma)$. From a similar analysis of the directional derivatives in directions orthogonal to δ_σ, it follows that the mixed cell is $\{f(\sigma) + sx \mid x \in \nu_\sigma\}$ is the image of $\overline{\mu_\sigma}$ onto $\mathrm{aff}(\nu_\sigma)$. Combining these two results, we obtain the following identity of the polyhedral cell:

$$\overline{\mu_\sigma} = (1-s) \cdot \delta_\sigma \oplus s \cdot \nu_\sigma,$$

which matches the definition of the mixed cell given in Section 2.7.4, Definition (5).

Parameterization of the weight function For ease of manipulating the piecewise quadratic weight functions, we define the weight function on a simplex by values at its vertices and at the midpoints of its edges. A quadratic function in d-dimensions, has $d(d+1)/2$ quadratic terms, d linear terms and a single constant. In total, there are $(d+1)(d+2)/2$ degrees of freedom. On the other hand, a d-dimensional simplex has $d+1$ vertices and $d(d+1)/2$ edges. If the vertices are affinely independent, the quadratic function is uniquely determined by specifying its value at the vertices on every vertex and at the midpoints of the edges.

To show that the weight function is continuous, let $\sigma, \sigma' \in T$ be two full dimensional cells and τ their common n-dimensional face. Restricted to τ, W^* is quadratic and interpolates the weights at the vertices and at the midpoints of the edges of τ. These weights determine the quadratic weight function uniquely and $W^*|_\tau = W_\sigma|_\tau = W_{\sigma'}|_\tau$.

4.5 Existing schemes

The new interpolation scheme, to be presented in Section 4.6, generalizes two existing methods. In this section, we describe these methods in terms of the definitions of the previous sections. An example in 2D of both methods is shown in Figure 4.6.

We first show how skin surfaces fit in our approach. Recall that, the radii of the input balls are shrunk in the construction of a skin surface.

In Chapter 3, we adapted the construction of skin surfaces such that the weight function interpolates the weights of the initial balls to the effect that the input balls are contained in the resulting skin surface. We call the skin surface defined in this way the extended skin surface.

Both methods define a one-parameter family of envelope surfaces, depending only on the shrink factor s. For a shrink factor equal to one, the weight functions of both methods are equal and therefore their envelope surfaces. We also use this weight function as starting point for the new interpolation scheme. This initial weight

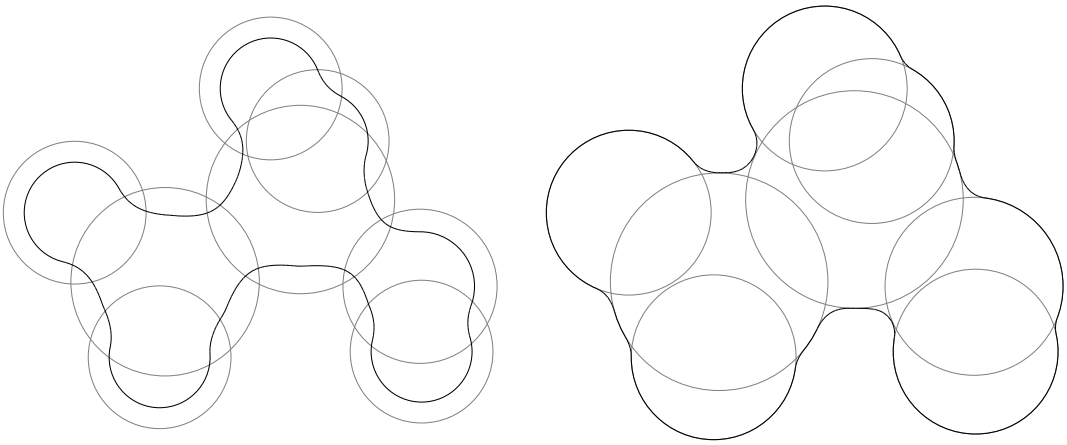

(a) (b)

Figure 4.6: Two existing interpolation schemes. On the left the skin curve [48] ($s = 1/2$) and on the right the extended skin curve presented in Chapter 3.

function, denoted by W^*, is presented in the next paragraph. The two schemes are then derived. In [48], Edelsbrunner proposes to deform skin surfaces by adding a scalar to all weights. This deformation works in general for envelope surfaces and is discussed at the end of the section.

Initial weight function Let T be the (weighted) Delaunay triangulation of the input set $\mathcal{P}$. The initial weight function W^* is piecewise quadratic with respect to T. Each quadratic function W^*_σ on a d-simplex σ is uniquely defined by the following two conditions:

- W^*_σ interpolates the weights on vertices of σ
- the defining matrix is the identity matrix.

We show that the W^* satisfies the Monotonous Transition Condition.

Lemma 36. *The weight function W^* satisfies the Monotonous Transition Condition.*

Proof. In view of Lemma 3, the graph of W^*_1, defined by $W^*_1(p) = \|p\|^2 - W^*(p)$, is the lower convex hull of the set of lifted points $\Pi(\mathcal{P})$, cf. Equation (2.2), which is convex. Note that W^*_1 is simplexwise affine.

Since W^*_1, is convex, the function W^*_1 satisfies the Monotonous Transition Condition by Lemma 69, and W^* satisfies the Monotonous Transition Condition as well. □

The weight function W^* is a good starting point for our construction since the envelope surface it defines is the boundary of the union of the input balls.

Lemma 37. *The envelope surface defined by the initial weight function W^* is equal to the boundary of the union of the input balls.*

Proof. The weight function interpolates the weights on centers of the input balls. Hence, the input balls are contained in the envelope surface. It remains to show that the approximating balls lie within the union of the input balls.

Let p' be a point in D and let $\sigma \in T$ be the simplex containing p'. Write p' as a convex combination of the vertices p_i of σ: $p' = \sum_i \gamma_i p_i$ with $\gamma_i \geq 0$ and $\sum_i \gamma_i = 1$. Since $W^*(x) = \|p\|^2 + W_1^*(p)$, so $\Phi(x,p) = \|x\|^2 - 2\langle p, x\rangle - W_1^*(p)$ is affine in p if p ranges over σ and

$$\Phi(x, \sum_i \gamma_i p_i) \quad = \quad \sum_i \gamma_i \Phi(x, p_i)$$

Let $x \in \mathbb{R}^d$ be a point outside the union of the balls in $\mathcal{P}$. Then, $\Phi(x, p_i)$ is positive, for any weighted point $(p_i, P_i) \in \mathcal{P}$, since $W^*(p_i) = P_i$. But then $\Phi(x, p') > 0$ and, hence, x also lies outside the ball centered at p'. □

Recall from the end of Section 4.4 that each quadric is uniquely determined by specifying the function values at vertices and at the midpoints of the edges. We denote the initial weight at the vertex $p_i \in \mathcal{P}$ of T with W_i^* and the initial weight at the midpoint of the edge $p_i p_j$ with W_{ij}^*. The initial weight function W^* interpolates the weights of the input balls if the value W_i^* is set to the weight of the input ball centered at p_i. Moreover, the defining matrix of $W^*|_\sigma$ is the identity matrix if the quadratic weight function restricted to any edge has leading coefficient equal to one. Let p_i, p_j be two vertices of σ then the quadratic weight function on the edge p_i, p_j has leading coefficient equal to one if the weight W_{ij}^* at the midpoint of the edge is:

$$W_{ij}^* = \frac{W_i^* + W_j^*}{2} - \frac{\|p_i - p_j\|^2}{4}.$$

Skin surfaces The weight function for a skin surface is obtained by multiplying the initial weight function with s:

$$W^s(x) = s\, W^*(x).$$

In terms of the parameterization of the weight function, the weight function is obtained by multiplying the weights at the vertices and at the midpoints of the edges with s.

Corollary 38. *For $s < 0 < 1$, the skin surface $\mathrm{skn}^s(\mathcal{X})$ is tangent continuous.*

Proof. The eigenvalues of the quadratic weight function restricted to a simplex are all equal to s. Hence, W^s satisfies the Strict Convexity Condition. By Lemma 36, W^* satisfies the Monotonous Transition Condition and so does W^s. Hence, the envelope surface obtained from the weight function W^s is C^1, which is also shown in [48]. □

Since all weights are multiplied by s, the weights on the vertices and edges controlling the weight functions are also multiplied by s. The weight of the input balls are scaled with a factor s, and therefore the input balls are not contained in the skin surface. See also Figure 4.6(a).

Extended skin surfaces The fact that the input balls are shrunk in the construction of skin surfaces makes them not directly suitable for the construction of envelope surfaces containing the input balls. In Chapter 3, we proposed a method that first multiplies the weights of the input balls with $1/s$, and then computes the skin surface with shrink factor s. An example of this construction is shown in Figure 4.6(b).

Since the extended skin surface is a skin surface (of the grown set of input balls), it is a C^1-manifold, as is shown in Corollary 38. If we combine the phases of growing the balls and multiplying the weight we obtain a deformation scheme of the weights on the vertices and edges. The weights on the vertices do not change. The weight at the midpoint of an edge p_ip_j is given by:

$$W_{ij} = \frac{W_i + W_j}{2} - s\frac{\|p_i - p_j\|^2}{4}.$$

During the growth of the input balls the (weighted) Delaunay triangulation may change, see Section 3.3. If we only increase the weights on the edges do not adapt the triangulation accordingly, the Monotonous Transition Condition is not satisfied and the envelope surface is not C^1. In Section 3.3, we developed an algorithm that maintains the Delaunay triangulation of the set of growing balls.

Deforming envelope surfaces A method to deform skin surfaces is to increase the weight by adding a scalar. The increased weight function also satisfies the Strict Convexity Condition and the Monotonous Transition Condition since the derivative of the weight function does not change. Therefore, the envelope surface is C^1. In fact, this surface is another level set of H. Since H is C^1, a level set of H is C^1 if it does not pass through a critical point of H. In view of Proposition 30, part 3 this is the case if the envelope surface does not intersect the zero set of the deformed weight function.

An example of a deforming skin curve obtained in this way is shown in Figure 4.8(e). For level sets of a curve constructed by the new interpolation scheme described in the next section, see Figure 4.8(f).

4.6 The new interpolation scheme

The schemes described in the previous section are global in the sense that they define a one parameter family of weight functions. In this section we propose an adaptive interpolation scheme. The advantage of this scheme is that it adapts the weight function locally allowing for a local control on the envelope surface. However,

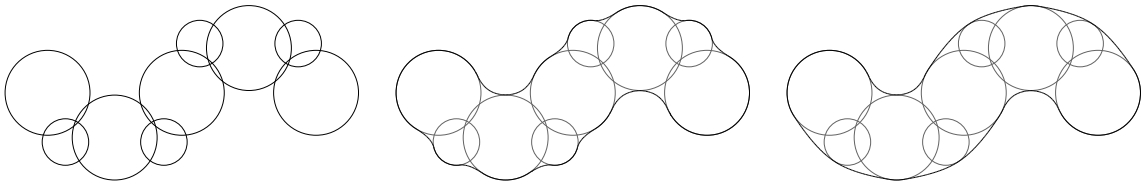

(a) Input set (b) Extended skin curve (c) The new interpolation scheme

Figure 4.7: The extended skin surface defines concave patches between the input balls. With the new interpolation scheme it is possible to interpolate both with convex and concave patches.

the envelope surface is only C^1 if Strict Convexity Condition and the Monotonous Transition Condition are satisfied, cf. Lemma 31.

Validation The Strict Convexity Condition is satisfied if the Hessian of the function $W_{\sigma,1} : D \to \mathbb{R}$, defined by $W_{\sigma,1}(p) = \|p\|^2 - W_\sigma(p)$, is positive definite, cf. Lemma 67. For a quadratic function, this is the case if the eigenvalues of the defining matrix are positive. Hence, the eigenvalues of the defining matrix of W_σ have to be smaller than one. The Strict Convexity Condition with respect to a simplex σ only depends on W_σ and therefore only on the weights at the vertices and at the midpoints of the edges of σ.

To see if the Monotonous Transition Condition is satisfied, let $\sigma, \sigma' \in T$ be two full dimensional cells with a common facet τ and let ν be the normal of τ directed from σ to σ'. Let the map w be defined on τ by $w(p) = W'_{\sigma'}(p;\nu) - W'_\sigma(p;\nu)$. Since the gradient of a quadratic function is a linear map, w is linear. Therefore, $w(p) \geq 0$ for all $p \in \tau$ iff $w(p) \geq 0$ for every vertex p of τ. Hence, the Monotonous Transition Condition is satisfied at every point of τ iff $w(p) \geq 0$ for every vertex p of τ. The Monotonous Transition Condition of a face τ incident to σ and σ' depends on W_σ and $W_{\sigma'}$ and therefore on the weights at the vertices and at the midpoints of the edges of σ and σ'.

Changing the weights Conceptually, the algorithm is similar to the extended skin surface algorithm. First construct the initial weight function W^* and then continuously increase the weights on the midpoints of the edges. We define a local growth factor t_{ij} for every edge p_ip_j and parameterize the weight at the midpoint of an edge by:

$$W_{ij} = \frac{W_i + W_j}{2} + (t_{ij} - 1)\frac{\|p_i - p_j\|^2}{4}.$$

Note that the envelope surface is an extended skin surface with $s = 1 - t_{ij}$, if all growth factors are equal. Initially all scalars t_{ij} are equal to zero. We continuously increase the local growth factors as long as the Strict Convexity Condition and the

Monotonous Transition Condition are satisfied. The envelope surface is therefore C^1. Since the conditions are locally determined, we can fix the local growth factors that would invalidate one of the conditions and increase the other local growth factors.

Another method fixes the local growth factor on edges between disjoint balls in an early stage and increases the local growth factor at midpoints of the edges between intersecting weighted points as much as possible. This gave us the interpolation on the fingers of the hand in Figure 4.8.

Examples Envelope curves that have an increasing weight on the midpoint of an edge are shown in Figure 4.1. Two examples of the interpolation scheme of [74] versus the local interpolation scheme are given in Figure 4.7 and 4.8. By construction, the extended skin surface interpolates the input balls with concave patches. The local interpolation scheme, on the other hand, allows for envelope surfaces with interpolating patches that are both convex and concave, viz. Figure 4.7. In Figure 4.8, we increased the weight on the edges in the direction of the fingers and not on edges between balls of different fingers. The result is an envelope surface that interpolates nicely in the direction of the fingers. The extended skin curve is either bumpy (b), or contains patches in between fingers (c).

The envelope surfaces in Figure 4.9 show the flexibility of these surfaces for only four input balls Figure 4.9(a). The envelope can be decomposed into quadric patches determined by simplices of the triangulation, see Section 4.4. The patches are color-coded by the dimension of this simplex. Figure 4.9(b)—(d) show envelope surfaces for which the local growth factor at the midpoints of the edges is increased with the same amount. These envelope surfaces are also extended skin surfaces. In Figure 4.9(e)—(g) we increased the local growth factor on the midpoints of one, two and three edges, respectively. The spheres are connected due to this increase.

4.7 Conclusions and future work

We constructed a class of smooth surfaces as the envelope of a set of balls. The theory behind envelope surfaces generalizes the concepts used for skin surfaces. The set of balls is defined by a convex set of centers and a weight function assigning a squared radius to each center. If the weight function is piecewise quadratic, then the envelope surface is also piecewise quadratic. The envelope surface is shown to be C^1 if two conditions on the weight function hold. These conditions can be verified automatically.

We assume that there is a topological relation between the weight function W and the function H defining the envelope surface:

Conjecture 39. *The isosurface of W and $-H$ are isotopic for all isovalues.*

The conjecture is true for skin surfaces and if W is differentiable then the critical points of W and H coincide and the critical values are equal with opposite signs, cf. Proposition 30.

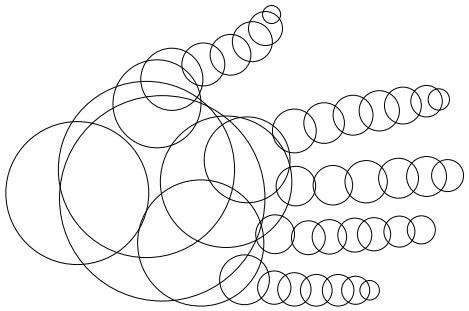

(a) An approximating set of circles

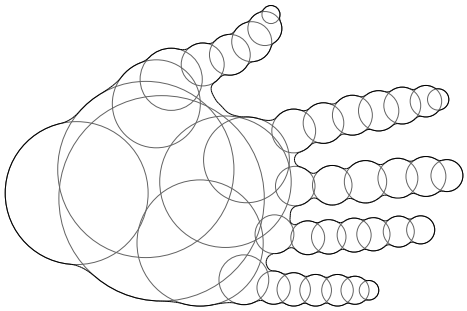

(b) The extended skin curve

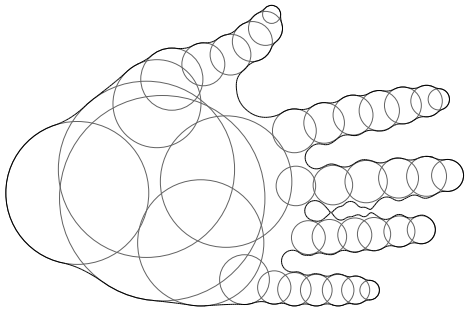

(c) The extended skin curve with a lower shrink factor

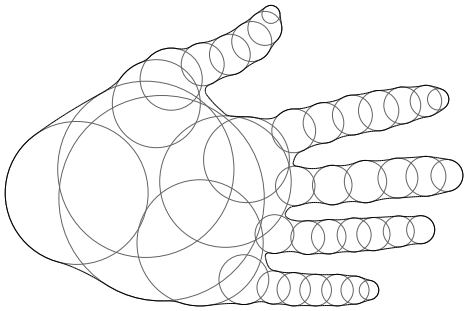

(d) The new interpolation scheme

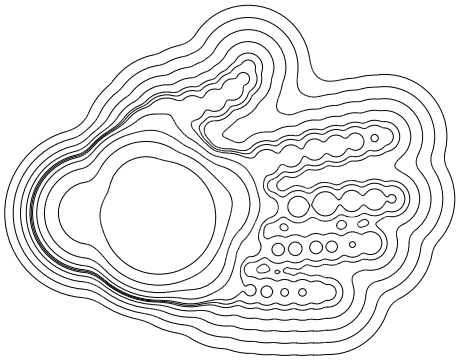

(e) Different level sets of H defined by the extended skin surface

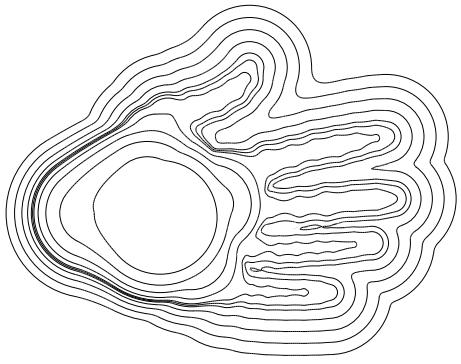

(f) Different level sets of H defined by the new interpolation scheme

Figure 4.8: Different interpolation schemes

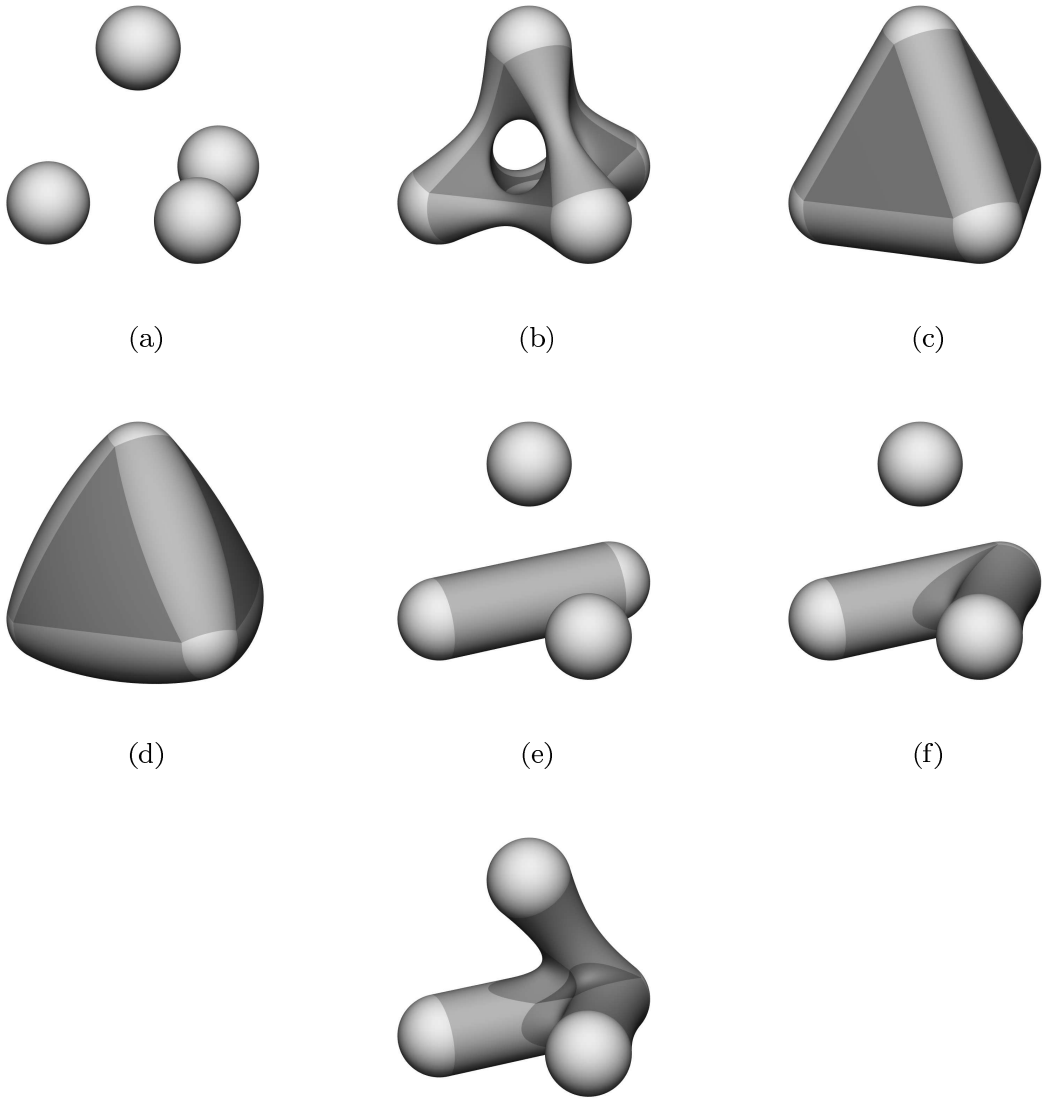

(a) (b) (c)

(d) (e) (f)

(g)

Figure 4.9: Envelope surfaces of four balls with different values for the local growth factors.

Chapter 5

Meshing Skin Surfaces

This chapter appeared as [75] and is accepted to Computational Geometry, Theory and Applications as [76].

5.1 Introduction

We present an algorithm for meshing skin surfaces with guaranteed topology. The algorithm presented in this chapter constructs a mesh isotopic to the skin surface in two steps: it constructs a coarse, isotopic mesh which is subsequently improved by refinement algorithms. The complexity of the coarse mesh is quadratic in the number of input balls, and is independent of the shrink factor. This is worst case optimal. For the second step a broad range of refinement algorithms can be used. Existing algorithms may have to be adapted slightly to ensure the isotopy. We show how this is done for the refinement algorithms of Chew [34] and Kobbelt [69]. The $\sqrt{3}$-subdivision algorithm by Kobbelt is very fast, and refines the size of the triangles. However, it does not improve the quality of the mesh elements in terms of angle size. Chew's algorithm improves the quality of the mesh in terms of the angles and the size of the triangles in terms of a measure like the circumradius or the length of the longest edge. The quality mesh is suitable for numerical computations. Our version of these algorithms preserve the isotopy property. Methods like the one presented in [44] are also well suited for mesh enhancements.

Related work Most existing algorithms for meshing implicit surfaces do not guarantee topological equivalence of the surface and the mesh constructed. The marching cubes algorithm [82] subdivides a region into cubes and triangulates the surface within these cubes based on whether the vertices of the cube lie inside or outside the cube. A variant of this algorithm that follows the surface is presented in [17].

The marching triangulation method [62] extends a small initial mesh by walking over the implicit surface, starting from a seed point. Our paper [73] presents a marching triangulation method for meshing skin surfaces by carefully choosing the step size during the walk over the mesh. However, as the shrink factor goes to one or to zero, the size of the mesh goes to infinity.

Several algorithms [27, 29, 30, 51] have been proposed to construct a topologically correct mesh approximating a skin surface in the special case of a shrink factor 1/2. It is likely that these algorithms can be generalized to work for arbitrary shrink factors, but this will result in a denser mesh in order to guarantee the topology and will therefore yield slower running times.

The algorithms in [27, 51] construct a dense point sample on the skin surface and use a three-dimensional Delaunay triangulation to extract the mesh from these points. These algorithms are rather slow compared to the other algorithms [29, 30].

Another approach is presented in [29]. The algorithm is competitive in terms of running times. We could not verify the claim that the mesh produced by the algorithm in this paper is homeomorphic to the skin surface. In fact, in [30] the authors say the following. "This method improved the efficiency dramatically (with respect to [27]) but resulted in robustness problem of the implementation ...The accumulation of numerical errors in the approximation of Morse-Smale complex led to inconsistent critical points after eliminating noisy critical points inaccurately."

The recent algorithm in [30] uses the theoretical results from [27] combined with ideas of the advancing front from [29]. The mesh is constructed by advancing the mesh over the skin surface while maintaining the restricted Delaunay triangulation.

Another method for visualizing molecules uses Molecular Surfaces [38]. Visualization algorithms for this type of surface are presented in [14, 15]. The algorithms presented in [21, 85] are the first general methods guaranteeing topological equivalence of the implicit surface and the mesh.

Contribution The approach to meshing skin surfaces described in this chapter is new. The main contribution compared to [27] is that our approach works for any shrink factor. We also establish isotopy, which is stronger than topological equivalence. Our algorithm is more flexible in the sense that we generate a coarse mesh that is isotopic to the skin surface that can subsequently be refined by various algorithms, as shown in Section 5.3.4, whereas the algorithm in [27] immediately constructs a homeomorphic quality mesh. Further, our algorithm is much faster. It constructs a mesh in minutes where the algorithm presented in [27] takes hours.

On the theoretical side, we analyze the structure of the mixed complex and decompose the mixed cells into tetrahedra. Within a tetrahedron the intersection with the skin surface is either empty or a topological disk. It is fairly easy to extract the isotopic mesh from this tetrahedral complex by a marching tetrahedra algorithm.

Outline In Section 5.2 we extend the theory of skin surfaces as presented in [48]. We start by introducing a hierarchical combinatorial structure on the mixed complex. With each face of this complex we associate an anchor point, which plays a crucial role in the meshing algorithm. Section 5.3 describes the construction of the coarse mesh and establishes the isotopy between this mesh and the skin surface. In Section 5.3.4, we describe two methods to improve the coarse mesh via (i) subdivision of the triangles and (ii) improvement of the quality of the triangles with regard to the size of the minimal angle. Finally, we describe our implementation and give experimental results in Section 5.4 and 5.5.

5.2 Definitions

This section first briefly reviews skin surfaces and then introduces some new concepts specific to the meshing algorithm. For a more thorough introduction to skin surfaces, we refer to Section 2.7 and to [48] where they were originally introduced.

5.2.1 Skin surfaces

A skin surface is defined in terms of a finite set of weighted points $\mathcal{P}$ and a shrink factor s, with $0 \leq s \leq 1$. A weighted point $\hat{p} = (p, P) \in \mathbb{R}^d \times \mathbb{R}$ corresponds to a ball with center p and radius $\sqrt{P}$. A pseudo-distance between two weighted points is given by:

$$\pi(\hat{p}, \hat{q}) = \|p - q\|^2 - P - Q,$$

where $\hat{p} = (p, P)$, $\hat{q} = (q, Q)$ and $\|\cdot\|$ denotes the Euclidean distance. The pseudo-distance $\pi(\hat{p}, x)$ of a weighted point $\hat{p}$ to an (unweighted) point x is the pseudo-distance of $\hat{p}$ to the weighted point centered at x with zero weight. Two weighted points with zero distance are called *orthogonal*. An *orthosphere* of a set of weighted points $\mathcal{P}$ is, by definition, a sphere orthogonal to each of the weighted points in $\mathcal{P}$.

The space of weighted points inherits a vector space structure from $\mathbb{R}^{d+1}$ via the bijective map $\Pi : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}^{d+1}$, defined by $\Pi(\hat{p}) = (x_1, \ldots, x_d, \|p\|^2 - P)$, with $p = (x_1, \ldots, x_d)$. Addition of two weighted points and the multiplication of a weighted point by a scalar are defined in the vector space structure inherited under Π. For further reading on the space of circles and spheres we refer to [84, 41].

Starting from a weighted point $\hat{p} = (p, P)$, the shrunk weighted point $\hat{p}^s$ is defined as $\hat{p}^s = (p, s \cdot P)$. The set $\mathcal{P}^s$ is the set obtained by shrinking every weighted point of $\mathcal{P}$ by a factor s.

The skin surface $\mathrm{skn}^s(\mathcal{P})$ and its body $\mathrm{bdy}^s(\mathcal{P})$ associated with a set of weighted points $\mathcal{P}$, are defined by

$$\mathrm{bdy}^s(\mathcal{P}) = \cup(\mathrm{conv}\,\mathcal{P})^s$$
$$\mathrm{skn}^s(\mathcal{P}) = \partial\ \mathrm{bdy}^s(\mathcal{P}).$$

Here $\mathrm{conv}(\mathcal{P}) \subset \mathbb{R}^d \times \mathbb{R}$ is the convex hull – with respect to the vector space structure inherited under Π – of a set of weighted points $\mathcal{P}$, whereas ∂ denotes the boundary – in $\mathbb{R}^d$ – of the union of the corresponding set of set of balls. For a skin curve in 2D associated with two weighted points: see Figure 5.1.

5.2.2 Delaunay triangulation

The Delaunay triangulation and Voronoi diagram are used to decompose the skin surface into patches of spheres and hyperboloids. We briefly give the definition of these structures and mention some properties.

The (weighted) Voronoi diagram (or: the power diagram) $\mathrm{Vor}(\mathcal{P})$ of a set of weighted points $\mathcal{P}$ is the subdivision of $\mathbb{R}^d$ into cells $\nu_{\mathcal{X}}$ that have smaller power

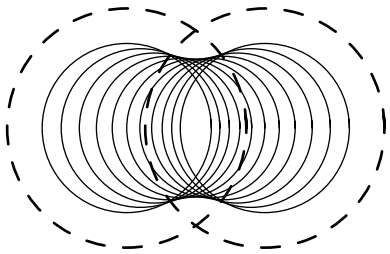

Figure 5.1: The skin curve of two weighted points (the two dashed circles). The smaller circles form a subset of the shrunk convex hull of the input points. Its boundary forms the skin curve.

distance to the weighted points in $\mathcal{X} \subseteq \mathcal{P}$ than to any other weighted point in $\mathcal{P}$:

$$\nu_{\mathcal{X}} = \bigcap_{\hat{p} \in \mathcal{X}, \hat{p}' \in \mathcal{P}} \{x \in \mathbb{R}^d | \pi(\hat{p}, x) \leq \pi(\hat{p}', x)\}.$$

Observation 40. *Let $y_{\hat{p},\hat{p}'}$ be a point with the same power distance to $\hat{p}$ and $\hat{p}'$, then $\nu_{\mathcal{X}} = \bigcap_{\hat{p} \in \mathcal{X}, \hat{p}' \in \mathcal{P}} \{x \in \mathbb{R}^d | \langle x - y_{\hat{p},\hat{p}'}, p' - p \rangle \leq 0\}$*

Proof. We have

$$\begin{aligned}
& \pi(\hat{p}, x) \leq \pi(\hat{p}', x) \\
\equiv \quad & \|p - x\|^2 - P \leq \|p' - x\|^2 - P' \\
\equiv \quad & 2\langle x, p' - p \rangle \leq P - P' + \|p'\|^2 - \|p\|^2 \\
\equiv \quad & 2\langle x, p' - p \rangle \leq 2\langle y_{\hat{p},\hat{p}'}, p' - p \rangle \\
\equiv \quad & \langle x - y_{\hat{p},\hat{p}'}, p' - p \rangle \leq 0.
\end{aligned}$$

Hence, the proof follows from the definition of a Voronoi cell. □

The dual of the Voronoi diagram is the Delaunay triangulation (or: regular triangulation) $\mathrm{Del}(\mathcal{P})$. We denote a Delaunay simplex of a set $\mathcal{X} \subseteq \mathcal{P}$, with $\nu_{\mathcal{X}} \neq \emptyset$, by $\delta_{\mathcal{X}}$. Recall that $\delta_{\mathcal{X}} = \mathrm{conv}(\{p | \hat{p} \in \mathcal{X}\})$. If $\mathcal{X} \subsetneq \mathcal{X}'$ and $\nu_{\mathcal{X}'} \neq \emptyset$, then $\nu_{\mathcal{X}'}$ is a proper face of $\nu_{\mathcal{X}}$ and $\delta_{\mathcal{X}}$ is a proper face of $\delta_{\mathcal{X}'}$.

The affine hulls of a Delaunay simplex $\delta_{\mathcal{X}}$ and its dual Voronoi cell $\nu_{\mathcal{X}}$ are complementary and orthogonal. Hence, the affine hulls of $\delta_{\mathcal{X}}$ and $\nu_{\mathcal{X}}$ always intersect in a single point, the focus $f(\mathcal{X})$ of $\mathcal{X}$.

General position In the remainder of this chapter we assume general position, by which we mean that no $d+2$ weighted points are equidistant to a point in $\mathbb{R}^d$ and no $k+2$ centers of weighted points lie on a common k-flat for $k = 0, \ldots, d-1$. Several methods like [50] exist to symbolically perturb a data set and ensure these conditions. Note that, under this genericity condition, an orthosphere of a set $\mathcal{X}$ only exists if $|\mathcal{X}| \leq d+1$.

Consider a Delaunay cell $\delta_{\mathcal{X}'}$ and one of its faces $\delta_{\mathcal{X}}$, with $\mathcal{X} \subsetneq \mathcal{X}'$. Their duals are respectively a face of a Voronoi cell and the Voronoi cell itself. There is

a halfspace through $\delta_{\mathcal{X}}$ containing $\delta_{\mathcal{X}'}$ and a halfspace through $\nu_{\mathcal{X}'}$ containing $\nu_{\mathcal{X}}$ such that their normals point in opposite directions.

Lemma 41. *Let $\delta_{\mathcal{X}}, \delta_{\mathcal{X}'} \in \mathrm{Del}(\mathcal{P})$, such that $\delta_{\mathcal{X}}$ is a proper face of $\delta_{\mathcal{X}'}$ and let $u = x'_\delta - x_\delta$ with $x_\delta \in \delta_{\mathcal{X}}$, $x'_\delta \in \mathrm{int}(\delta_{\mathcal{X}'})$. Then*

1. *$\langle u, x_\nu - x'_\nu \rangle = 0$, for $x_\nu, x'_\nu \in \nu_{\mathcal{X}'}$.*
2. *$\langle u, x_\nu - x'_\nu \rangle < 0$, for $x_\nu \in \nu_{\mathcal{X}} \setminus \nu_{\mathcal{X}'}$, $x'_\nu \in \nu_{\mathcal{X}'}$*

Proof. Claim (1) follows directly from the orthogonality of $\delta_{\mathcal{X}'}$ and $\nu_{\mathcal{X}'}$, cf. Observation 1. Hence, claim (2) is independent of the choice of x'_ν.

For the proof of claim (2), let $m = |\mathcal{X}|$, $n = |\mathcal{X}'|$ and $\mathcal{X}' = \{\hat{p}_1, \ldots, \hat{p}_n\}$, such that $\hat{p}_i \in \mathcal{X}$, for $i \leq m$. Write x_δ and x'_δ in barycentric coordinates: $x_\delta = \sum \gamma_i \cdot p_i$, $x'_\delta = \sum \gamma'_i \cdot p_i$ with $\sum \gamma_i = \sum \gamma'_i = 1$, $\gamma_i, \gamma'_i \geq 0$. Since $x_\delta \in \delta_{\mathcal{X}}$, $\gamma_i = 0$ for $i \geq m+1$, and $\gamma'_i > 0$ since $x'_\delta \in \mathrm{int}(\delta_{\mathcal{X}'})$. Rewrite u as:

$$\begin{aligned} u &= x'_\delta - x_\delta \\ &= \sum_{i=1}^{m} (\gamma'_i - \gamma_i) p_i + \sum_{i=m+1}^{n} \gamma'_i p_i \\ &= \sum_{i=1}^{m} (\gamma'_i - \gamma_i)(p_i - p_1) + \sum_{i=m+1}^{n} \gamma'_i (p_i - p_1) \end{aligned}$$

Expanding $\langle u, x_\nu - x'_\nu \rangle$ yields:

$$\begin{aligned} \langle u, x_\nu - x'_\nu \rangle &= \sum_{i=1}^{m} (\gamma'_i - \gamma_i) \langle p_i - p_1, x_\nu - x'_\nu \rangle \\ &\quad + \sum_{i=m+1}^{n} \gamma'_i \langle p_i - p_1, x_\nu - x'_\nu \rangle \end{aligned}$$

From Observation 40, with $y_{\hat{p}_i, \hat{p}_1} = x'_\nu$, it follows that $\langle p_i - p_1, x_\nu - x'_\nu \rangle$ is not positive. Moreover, the inner product is zero if and only if $\hat{p}_i \in \mathcal{X}$. Hence, the elements of the first sum are zero and the elements of the second sum are negative, so $\langle u, x_\nu - x'_\nu \rangle < 0$. Note that for $\hat{p}_1$ we can substitute any weighted point in $\mathcal{X}$. □

Although the meshing algorithm generalizes to any dimension, the main application is in $\mathbb{R}^3$. Therefore we present the algorithm in three space.

5.2.3 The mixed complex

We already described the mixed complex in Section 2.7.4. In this section we give a more thorough presentation of the mixed complex and derive its polyhedral structure.

The mixed complex $\mathrm{Mix}^s(\mathcal{P})$, associated with a scalar $s \in [0, 1]$, is an intermediate complex between the Delaunay triangulation and the Voronoi diagram. Each mixed cell in the mixed complex is obtained by taking Minkowski sums of shrunk Delaunay simplices and their dual Voronoi cells.

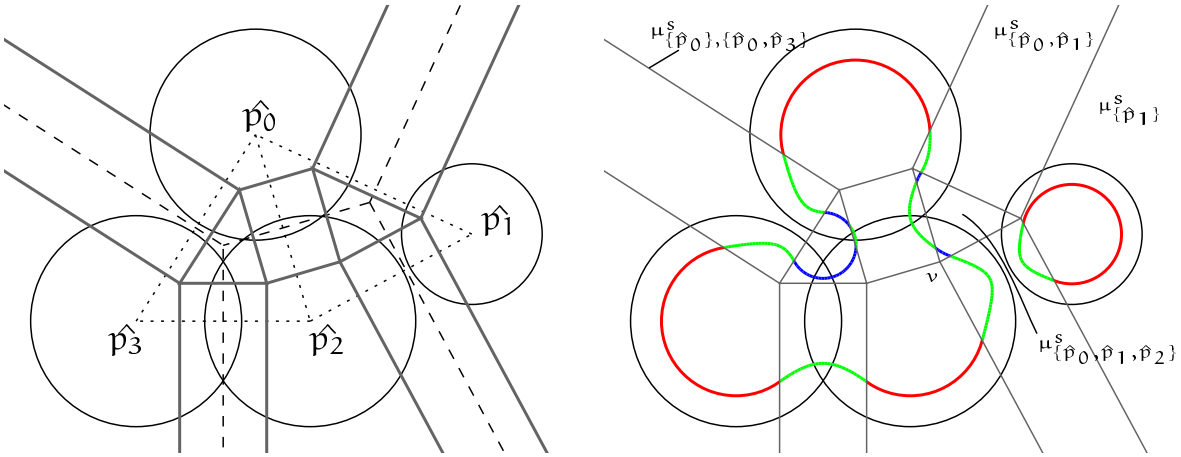

(a) The Delaunay triangulation (dashed), the Voronoi diagram (dotted) and the Mixed complex (solid).

(b) The mixed complex and some of the labels. Note that $\nu = \mu^s_{\{\hat{p}_0,\hat{p}_2\},\{\hat{p}_1,\hat{p}_2\}} = \mu^s_{\{\hat{p}_2\},\{\hat{p}_0,\hat{p}_1,\hat{p}_2\}}$.

Figure 5.2: The skin curve of four weighted points (the circles). Each mixed cell contains parts of an hyperbola or a circle.

Definition 42. *For* $\delta_{\mathcal{X}} \in \mathrm{Del}(\mathcal{P})$ *the* mixed cell $\mu^s_{\mathcal{X}}$ *is defined by* $\mu^s_{\mathcal{X}} = (1-s) \cdot \delta_{\mathcal{X}} \oplus s \cdot \nu_{\mathcal{X}}$.

Here $\cdot$ denotes the multiplication of a set by a scalar and $\oplus$ denotes the Minkowski sum. For $s = 0$ the mixed cell is the Delaunay cell. If s increases it deforms affinely into the Voronoi cell for $s = 1$.

Each mixed cell is a convex polyhedron since it is the Minkowski sum of two convex polyhedra. Based on the dimension of the Delaunay simplex, there are four types of mixed cells. A mixed cell of type ℓ corresponds to a Delaunay ℓ-cell and is of the form $\mu^s_{\mathcal{X}}$ with $|\mathcal{X}| = \ell + 1$. In 3D, mixed cells of type 3 are tetrahedra (shrunk Delaunay 3-cells) and mixed cells of type 0 are shrunk Voronoi 3-cells. A mixed cell of type 1 or 2 is a prism with respectively the shrunk Voronoi facet or the shrunk Delaunay facet as its base.

The intersection of the skin surface and a mixed cell is a piece of a sphere or a hyperboloid. In the plane, the intersection of a skin curve with a mixed cell is either part of a circle or hyperbola. An example of the mixed complex and a skin curve is given in Figure 5.2. All rectangles are mixed cells of type 1 and contain hyperbolic patches. The other cells contain circular arcs. Depending on whether the mixed cell is of type 0 or 2, the interior of the skin curve lies inside or outside the circle.

Within a mixed ℓ-cell $\mu^s_{\mathcal{X}}$, the skin surface is a quadratic surface of the form $I^{-1}_{\mathcal{X}}(0)$, where:

$$I_{\mathcal{X}}(x) = -\frac{1}{1-s} \sum_{i=1}^{\ell} x_i^2 + \frac{1}{s} \sum_{i=\ell+1}^{3} x_i^2 - R^2, \tag{5.1}$$

with $x = (x_1, x_2, x_3)$ and R^2 the weight of the weighted point in $\mathrm{aff}(\delta_{\mathcal{X}})$ centered at $f(\mathcal{X})$. More precisely, $\mathrm{skn}^s(\mathcal{X}) \cap \mu^s_{\mathcal{X}} = I^{-1}_{\mathcal{X}}(0) \cap \mu^s_{\mathcal{X}}$. The coordinate system is orthonormal with its origin at the focus of $\mathcal{X}$, and such that the first ℓ coordinates span the affine hull $\delta_{\mathcal{X}}$, see [48].

The following observation holds trivially for mixed cells of type 0 and 3. For mixed cells of type 1 and 2, the symmetry sets of the hyperboloids are the affine hulls of the corresponding Delaunay simplex and Voronoi cell. Hence it follows from the construction of the mixed cells.

Observation 43. *Each proper face of a mixed cell $\mu^s_{\mathcal{X}}$ is perpendicular to a symmetry set of $I_{\mathcal{X}}$*

Since the symmetry axis and the symmetry plane of the hyperboloid are perpendicular, each face of a mixed cell of type 1 or 2 is parallel to the other symmetry set.

Polyhedral complex The mixed complex is a polyhedral complex. The 3-cells of this polyhedral complex are formed by the mixed cells. We give a more detailed description of its structure.

Definition 44. *For $\mathcal{X}, \mathcal{X}' \subseteq \mathcal{P}$, with $\nu_{\mathcal{X}}, \nu_{\mathcal{X}'} \neq \emptyset$, a polyhedral cell $\mu^s_{\mathcal{X},\mathcal{X}'}$ is defined as $\mu^s_{\mathcal{X},\mathcal{X}'} = \mu^s_{\mathcal{X}} \cap \mu^s_{\mathcal{X}'}$.*

Edelsbrunner gives an intuitive picture of the mixed complex in [48]. Take the interval of d-dimensional affine subspaces of $\mathbb{R}^{d+1}$ defined by $x_{d+1} = s$, for $s \in [0, 1]$. Draw $\mathrm{Del}(P)$ in $x_{d+1} = 0$ and $\mathrm{Vor}(\mathcal{P})$ in $x_{d+1} = 1$. For each Delaunay simplex and corresponding Voronoi cell construct

$$\mu_{\mathcal{X}} = \mathrm{conv}(\delta_{\mathcal{X}} \cup \nu_{\mathcal{X}}).$$

All $\mu_{\mathcal{X}}$ are convex polyhedra of dimension $d+1$, their interiors mutually disjoint, and they decompose the strip between $x_{d+1} = 0$ and $x_{d+1} = 1$. The subspace $x_{d+1} = s$ intersects $\mu_{\mathcal{X}}$ in the mixed cell $\mu^s_{\mathcal{X}}$.

It is clear that a polyhedral cell $\mu^s_{\mathcal{X},\mathcal{X}'}$ is non-empty, for $0 < s < 1$, if the Delaunay and Voronoi cells of $\mathcal{X}$ and $\mathcal{X}'$ have a non-empty intersection. Or, equivalently, if $\nu_{\mathcal{X}\cap\mathcal{X}'}, \nu_{\mathcal{X}\cup\mathcal{X}'} \in \mathrm{Vor}(\mathcal{P})$. It is not enough for one of the two simplices to exists. E.g., let $\hat{p}_1, \hat{p}_2, \hat{p}_3$ be weighted points, the centers of which are the vertices of a triangle in a two-dimensional Delaunay triangulation and $\mathcal{X} = \{\hat{p}_1, \hat{p}_2\}$, $\mathcal{X}' = \{\hat{p}_3\}$. Then $\delta_{\mathcal{X}\cup\mathcal{X}'} \in \mathrm{Del}(\mathcal{P})$, but $\mathcal{X} \cap \mathcal{X}' = \emptyset$, hence $\delta_{\mathcal{X}\cap\mathcal{X}'} \notin \mathrm{Del}(\mathcal{P})$. On the other hand, let $\delta_{\mathcal{X}}$, $\delta_{\mathcal{X}'}$ be two Delaunay edges, that share a common vertex, but do not have an incident triangle in common, then $\delta_{\mathcal{X}\cap\mathcal{X}'} \in \mathrm{Del}(\mathcal{P})$, but $\delta_{\mathcal{X}\cup\mathcal{X}'} \notin \mathrm{Del}(\mathcal{P})$.

The following lemma describes the structure of the mixed complex.

Lemma 45. *A mixed cell $\mu^s_{\mathcal{X},\mathcal{X}'}$ is not empty iff $\nu_{\mathcal{X}\cap\mathcal{X}'}$ and $\nu_{\mathcal{X}\cup\mathcal{X}'}$ are nonempty In that case,*

$$\mu^s_{\mathcal{X},\mathcal{X}'} = (1 - s) \cdot \delta_{\mathcal{X}\cap\mathcal{X}'} \oplus s \cdot \nu_{\mathcal{X}\cup\mathcal{X}'}.$$

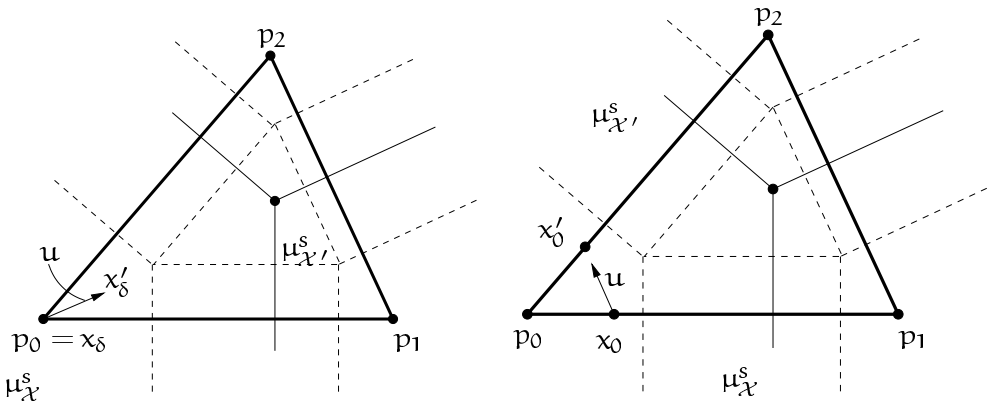

(a) $\mathcal{X} = \{\hat{p}_0\}, \mathcal{X}' = \{\hat{p}_0, \hat{p}_1, \hat{p}_2\}$, hence $\mathcal{X} \subseteq \mathcal{X}'$

(b) $\mathcal{X} = \{\hat{p}_0, \hat{p}_1\}, \mathcal{X}' = \{\hat{p}_0, \hat{p}_2\}$, hence $\mathcal{X} \not\subset \mathcal{X}'$ and $\mathcal{X}' \not\subset \mathcal{X}$

Figure 5.3: Illustration of the proof of Lemma 45.

Before we prove this lemma we first make some general remarks. Note that the lemma also holds if $\nu_{\mathcal{X}\cap\mathcal{X}'} = \emptyset$ or $\nu_{\mathcal{X}\cup\mathcal{X}'} = \emptyset$, since then $\mu^s_{\mathcal{X}} = \emptyset$.

Corollary 46. *If $\nu_{\mathcal{X}\cap\mathcal{X}'}$ and $\nu_{\mathcal{X}\cup\mathcal{X}'}$ are nonempty, then $\mu^s_{\mathcal{X},\mathcal{X}'} = \mu^s_{\mathcal{X}\cap\mathcal{X}',\mathcal{X}\cup\mathcal{X}'}$*

The corollary holds since, $\mu^s_{\mathcal{X}\cap\mathcal{X}',\mathcal{X}\cup\mathcal{X}'} = (1-s)\cdot\delta_{\mathcal{X}\cap\mathcal{X}'} \oplus s\cdot\nu_{\mathcal{X}\cup\mathcal{X}'}$. Hence, each non-empty polyhedral cell $\mu^s_{\mathcal{X},\mathcal{X}'}$ has a unique label, if $\mathcal{X} \subset \mathcal{X}'$. To gain some intuition for the lemma, take s equal to zero. Then the mixed complex is the Delaunay triangulation and indeed $\mu^0_{\mathcal{X},\mathcal{X}'} = \delta_{\mathcal{X}} \cap \delta_{\mathcal{X}'} = \delta_{\mathcal{X}\cap\mathcal{X}'}$. Conversely, for $s = 1$, the mixed complex is the Voronoi diagram and $\mu^1_{\mathcal{X},\mathcal{X}'} = \nu_{\mathcal{X}} \cap \nu_{\mathcal{X}'} = \nu_{\mathcal{X}\cup\mathcal{X}'}$.

Proof of Lemma 45. The proof is trivial if $\mathcal{X} = \mathcal{X}'$, hence we assume that $\mathcal{X} \neq \mathcal{X}'$. For simplicity, let $F = (1-s)\cdot\delta_{\mathcal{X}\cap\mathcal{X}'} \oplus s\cdot\nu_{\mathcal{X}\cup\mathcal{X}'}$.

From Definition (42) it follows that $F \subseteq \mu^s_{\mathcal{X}}$ and $F \subseteq \mu^s_{\mathcal{X}'}$, since $\delta_{\mathcal{X}\cap\mathcal{X}'} \subseteq \delta_{\mathcal{X}}, \delta_{\mathcal{X}'}$ and $\nu_{\mathcal{X}\cup\mathcal{X}'} \subseteq \nu_{\mathcal{X}}, \nu_{\mathcal{X}'}$. Hence, $F \subseteq \mu^s_{\mathcal{X},\mathcal{X}'}$.

For the opposite inclusion, we show that the two mixed cells lie in opposite halfspaces and intersect the bounding plane in F.

We distinguish two cases. First, consider the case where $\mathcal{X} \subseteq \mathcal{X}'$ or $\mathcal{X}' \subseteq \mathcal{X}$; See Figure 5.3(a). Without loss of generality we assume that $\mathcal{X} \subseteq \mathcal{X}'$. Let u be a vector perpendicular to $\delta_{\mathcal{X}}$ pointing from a point in $\delta_{\mathcal{X}}$ towards a point in the interior of $\delta_{\mathcal{X}'}$, such that $\langle u, x'_\delta - x_\delta\rangle > 0$, for $x_\delta \in \delta_{\mathcal{X}}$, $x'_\delta \in \delta_{\mathcal{X}'} \setminus \delta_{\mathcal{X}}$. Such a vector u exists, since $\delta_{\mathcal{X}}$ is a proper face of the convex polyhedron $\delta_{\mathcal{X}'}$. Note that u is perpendicular to $\delta_{\mathcal{X}}$. Lemma 41(2) states that $\langle u, x'_\nu - x_\nu\rangle < 0$, for $x_\nu \in \nu_{\mathcal{X}} \setminus \nu_{\mathcal{X}'}$, $x_\nu \in \nu_{\mathcal{X}'}$.

For each point x in a mixed cell $\mu^s_{\mathcal{X}}$ there exists a unique combination $x_\delta \in \delta_{\mathcal{X}}$, $x_\nu \in \nu_{\mathcal{X}}$, such that $x = (1-s)\cdot\delta_{\mathcal{X}} + s\cdot\nu_{\mathcal{X}}$, since $\delta_{\mathcal{X}}$ and $\nu_{\mathcal{X}}$ are affinely independent. Hence, since $F \subseteq \mu^s_{\mathcal{X}}$, a point $y_0 \in F$ can be uniquely written as $y_0 = (1-s)\cdot y^0_\delta + s\cdot y^0_\nu$ with $y^0_\delta \in \delta_{\mathcal{X}}$, $y^0_\nu \in \nu_{\mathcal{X}}$.

We analyze the sign of the inner product $\langle u, y - y_0\rangle$ for y subsequently in $\mu^s_{\mathcal{X}} \setminus F$, F and $\mu^s_{\mathcal{X}'} \setminus F$.

First, let $y \in \mu^s_{\mathcal{X}} \setminus F$. We write $y = (1-s)\cdot y_\delta + s \cdot y_\nu$, with $y_\delta \in \delta_{\mathcal{X}}$ and $y_\nu \in \nu_{\mathcal{X}} \setminus \nu_{\mathcal{X}'}$. The inner product $\langle u, y_\delta - y^0_\delta\rangle$ is zero since $y_\delta, y^0_\delta \in \delta_{\mathcal{X}}$ and $\langle u, y_\nu - y^0_\nu\rangle < 0$ by Lemma 41(2). Hence $\langle u, y - y_0\rangle < 0$ for $y \in \mu^s_{\mathcal{X}} \setminus F$.

Now assume that $y \in F$. Similar to y_0, we write $y = (1-s)\cdot y_\delta + s \cdot y_\nu$, with $y_\delta \in \delta_{\mathcal{X}}$, $y_\nu \in \nu_{\mathcal{X}'}$. The inner product $\langle u, y_\delta - y^0_\delta\rangle$ is zero since $y_\delta, y^0_\delta \in \delta_{\mathcal{X}}$ and $\langle u, y_\nu - y^0_\nu\rangle = 0$ by Lemma 41*(1)*.

Finally, assume that $y \in \mu^s_{\mathcal{X}'} \setminus F$, then we write y as $y = (1-s)\cdot y_\delta + s \cdot y_\nu$, with $y_\delta \in \delta_{\mathcal{X}'} \setminus \delta_{\mathcal{X}}$ and $y_\nu \in \nu_{\mathcal{X}'}$. The inner product $\langle u, y_\delta - y^0_\delta\rangle$ is positive by construction of u, and $\langle u, y_\nu - y^0_\nu\rangle = 0$, again by Lemma 41*(1)*. Hence $\langle u, y - y_0\rangle > 0$ for $y \in \mu^s_{\mathcal{X}'} \setminus F$.

Summarizing, we have:

$$\langle u, y - y_0\rangle \begin{cases} < 0, & \text{for } y \in \mu^s_{\mathcal{X}} \setminus F, \\ = 0, & \text{for } y \in F, \\ > 0, & \text{for } y \in \mu^s_{\mathcal{X}'} \setminus F. \end{cases}$$

Hence, $\mu^s_{\mathcal{X}}$ and $\mu^s_{\mathcal{X}'}$ lie in opposite halfspaces and meet only in F.

We continue with the proof of the second case. Assume that $\delta_{\mathcal{X}}$ is not a face of $\delta_{\mathcal{X}'}$ and vice versa. Then $\mathcal{X} \cap \mathcal{X}' \subsetneq \mathcal{X}, \mathcal{X}' \subsetneq \mathcal{X} \cup \mathcal{X}'$, viz. Figure 5.3(b). For this case the proof is similar, except for the construction of the vector u.

Let $x_I \in \delta_{\mathcal{X}\cap\mathcal{X}'}$. The Delaunay simplex $\delta_{\mathcal{X}\cap\mathcal{X}'}$ has at least co-dimension 2, since $|\mathcal{X} \cup \mathcal{X}'| - |\mathcal{X} \cap \mathcal{X}'| \geq 2$. Hence, the set of points orthogonal to $\delta_{\mathcal{X}\cap\mathcal{X}'}$ through x_I is at least 2-dimensional. We intersect this orthogonal set with a small sphere centered at x_I. If the radius is small enough, the intersection contains a point $x_0 \in \mathrm{int}(\delta_{\mathcal{X}})$ and $x'_0 \in \mathrm{int}(\delta_{\mathcal{X}'})$.

Let $u = \gamma \cdot (x'_0 - x_0)$, for some $0 < \gamma < 1$. By construction the triangle x_I, x_0, x'_0 is perpendicular to $\delta_{\mathcal{X}\cap\mathcal{X}'}$. Since $\|x_I - x_0\| = \|x_I - x'_0\|$, the triangle x_I, x_0, x'_0 is an isosceles triangle. Hence, the angles $\angle x_I, x_0, x'_0$, $\angle x_I, x'_0, x_0$ are equal and acute. As a result, for $y_0 \in \delta_{\mathcal{X}\cap\mathcal{X}'}$ we have:

$$\langle u, y - y_0\rangle \begin{cases} < 0, & \text{for } y \in \delta_{\mathcal{X}} \setminus \delta_{\mathcal{X}\cap\mathcal{X}'}, \\ = 0, & \text{for } y \in \delta_{\mathcal{X}\cap\mathcal{X}'}, \\ > 0, & \text{for } y \in \delta_{\mathcal{X}'} \setminus \delta_{\mathcal{X}\cap\mathcal{X}'}. \end{cases}$$

Note that u points from x_0 towards the interior of $\delta_{\mathcal{X}\cup\mathcal{X}'}$. Hence, u satisfies Lemma 41 with respect to $\delta_{\mathcal{X}}$ and $\delta_{\mathcal{X}\cup\mathcal{X}'}$. Using a similar argument, $-u$ satisfies Lemma 41 with respect to $\delta_{\mathcal{X}'}$ and $\delta_{\mathcal{X}\cup\mathcal{X}'}$. So, for $y_0 \in \nu_{\mathcal{X}\cup\mathcal{X}'}$ we have:

$$\langle u, y - y_0\rangle \begin{cases} < 0, & \text{for } y \in \nu_{\mathcal{X}} \setminus \nu_{\mathcal{X}\cup\mathcal{X}'}, \\ = 0, & \text{for } y \in \nu_{\mathcal{X}\cup\mathcal{X}'}, \\ > 0, & \text{for } y \in \nu_{\mathcal{X}'} \setminus \nu_{\mathcal{X}\cup\mathcal{X}'}. \end{cases}$$

Now we combine the results for the Delaunay simplices and the Voronoi cells to a statement for the mixed cell. Let $y_0 \in F$ and write y_0 as $y_0 = (1-s)\cdot y^0_\delta + s \cdot y^0_\nu$ with $y^0_\delta \in \delta_{\mathcal{X}}$, $y^0_\nu \in \nu_{\mathcal{X}'}$.

Write $y \in \mu^s_{\mathcal{X}} \setminus F$ uniquely as $y = (1-s)\cdot y_\delta + s \cdot y_\nu$, with $y_\delta \in \delta_{\mathcal{X}}$ and $y_\nu \in \nu_{\mathcal{X}}$. Since $y \notin F$, either $y_\delta \notin \delta_{\mathcal{X}\cap\mathcal{X}'}$ or $y_\nu \notin \nu_{\mathcal{X}\cup\mathcal{X}'}$. Expand the inner product

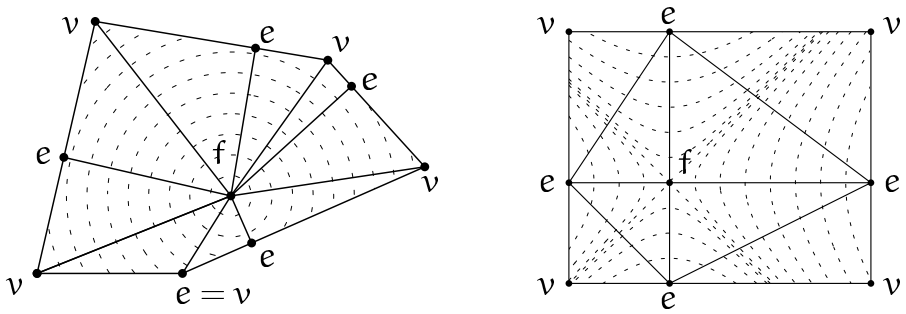

(a) Circular facets (b) Hyperbolic facets

Figure 5.4: The anchor points of two-dimensional polyhedral cells. Each anchor point is labeled by the type its cell (f for face, e for edge and v for vertex). The triangulation constructed in Section 5.3.2 is also shown.

$\langle u, y - y_0 \rangle$ to $(1-s) \cdot \langle u, y_\delta - y^0_\delta \rangle + s \cdot \langle u, y_\nu - y^0_\nu \rangle$. Using the estimates above, we obtain $\langle u, y - y_0 \rangle < 0$

A similar reasoning yields $\langle u, y - y_0 \rangle = 0$ for $y \in F$ and $\langle u, y - y_0 \rangle > 0$ for $y \in \mu^s_{\mathcal{X}} \setminus F$. □

Denote with $\mathrm{aff}(X)$ the affine hull of a set X.

Lemma 47. *For $\mathcal{X} \subset \mathcal{X}'$, $\mathrm{aff}(\mu^s_{\mathcal{X},\mathcal{X}'})$ and $\mathrm{aff}(\delta_{\mathcal{X}'}) \cap \mathrm{aff}(\nu_{\mathcal{X}})$ are complementary and orthogonal.*

Proof. To shorten notation we write $F = \mathrm{aff}(\mu^s_{\mathcal{X},\mathcal{X}'})$ and $G = \mathrm{aff}(\delta_{\mathcal{X}'}) \cap \mathrm{aff}(\nu_{\mathcal{X}})$. Recall from Lemma 45 that $\mu^s_{\mathcal{X},\mathcal{X}'} = (1-s) \cdot \delta_{\mathcal{X}} \oplus s \cdot \nu_{\mathcal{X}'}$.

The cells $\delta_{\mathcal{X}}$ and $\nu_{\mathcal{X}'}$ are affinely independent, hence $\dim F = \dim \delta_{\mathcal{X}} + \dim \nu_{\mathcal{X}'} = d + |\mathcal{X}| - |\mathcal{X}'|$. Further, $\delta_{\mathcal{X}'}$ and $\nu_{\mathcal{X}}$ are orthogonal and $\dim G = \dim \mathrm{aff}(\delta_{\mathcal{X}'}) - \dim \mathrm{aff}(\delta_{\mathcal{X}}) = |\mathcal{X}'| - |\mathcal{X}|$. Hence, the dimensions of F and G add up to d. Both $\delta_{\mathcal{X}}$ and $\nu_{\mathcal{X}'}$ are orthogonal to G, which shows the orthogonality of F and G. □

Corollary 48. *The dimension of a non-empty mixed cell $\mu^s_{\mathcal{X},\mathcal{X}'}$ in $\mathbb{R}^d$, is $d - |\mathcal{X} \cup \mathcal{X}'| + |\mathcal{X} \cap \mathcal{X}'|$.*

5.2.4 The anchor point

For the construction of the mesh we use the anchor point of a polyhedron.

Definition 49. *Let A be a convex set and p a point in $\mathbb{R}^3$. Then the* anchor point $a_p(A)$ *is the point in A closest to p.*

We are interested in the case where A is a polyhedral cell $\mu^s_{\mathcal{X},\mathcal{X}'}$, a Delaunay cell $\delta_{\mathcal{X}}$ or a Voronoi cell $\nu_{\mathcal{X}}$ and p is the focus $f(\mathcal{X})$. In fact, we use the anchor points of the polyhedral cells as vertices of a tetrahedral complex that decomposes the skin surface into topological disks.

We distinguish two types of critical points on a mixed cell $\mu^s_{\mathcal{X}}$, interior critical points are critical points of $I_{\mathcal{X}}$ contained in the interior of $\mu^s_{\mathcal{X}}$ and boundary critical points are critical points of $I_{\mathcal{X}}$ restricted to the boundary of $\mu^s_{\mathcal{X}}$. All critical points are anchor points of a face of the mixed cell, viz. Figure 5.4. However, not all anchor points are critical points, e.g. the point that is both the anchor point of a vertex and an edge in Figure 5.4(a).

Lemma 50. *A (boundary or regular) critical point of* $I_{\mathcal{X}}$ *on a polyhedral cell* $\mu^s_{\mathcal{X},\mathcal{X}'}$ *is the anchor point of* $\mu^s_{\mathcal{X},\mathcal{X}'}$ *or the anchor point of one of its faces with respect to* $f(\mathcal{X})$.

Proof. The focus $f(\mathcal{X})$ is the only critical point of the quadratic function $I_{\mathcal{X}}$. If $f(\mathcal{X})$ is contained in $\mu^s_{\mathcal{X},\mathcal{X}'}$, then it is the anchor point $a_{f(\mathcal{X})}(\mu^s_{\mathcal{X},\mathcal{X}'})$.

It remains to show that all boundary critical points are also anchor points. By Observation 43, a face of $\mu^s_{\mathcal{X},\mathcal{X}'}$ is either parallel or perpendicular to the symmetry sets of $I_{\mathcal{X}}$. Hence, if $f(\mathcal{X})$ projects onto the facet, then the facet has a boundary critical point. By definition, this point is the anchor point of the facet with respect to $f(\mathcal{X})$. □

Let $\mu^s_{\mathcal{X},\mathcal{X}'}$ be a common face of $\mu^s_{\mathcal{X}'}$ and $\mu^s_{\mathcal{X}}$. The following lemma shows that the anchor points of $\mu^s_{\mathcal{X},\mathcal{X}'}$ with respect to $f(\mathcal{X})$ and $f(\mathcal{X}')$ are identical.

Lemma 51. $a_{f(\mathcal{X})}(\mu^s_{\mathcal{X},\mathcal{X}'}) = a_{f(\mathcal{X}')}(\mu^s_{\mathcal{X},\mathcal{X}'})$.

Proof. Both $f(\mathcal{X})$ and $f(\mathcal{X}')$ lie on $\operatorname{aff}(\delta_{\mathcal{X}\cup\mathcal{X}'})$ and $\operatorname{aff}(\nu_{\mathcal{X}\cap\mathcal{X}'})$ by definition. Hence they lie on $\operatorname{aff}(\delta_{\mathcal{X}\cup\mathcal{X}'})\cap\operatorname{aff}(\nu_{\mathcal{X}\cap\mathcal{X}'})$, which is orthogonal to $\mu^s_{\mathcal{X},\mathcal{X}'}$ by Lemma 47. □

The following lemma gives a relation between the anchor point of a polyhedral cell and the anchor points of Delaunay simplices and Voronoi cells.

Lemma 52. $a_{f(\mathcal{X})}(\mu^s_{\mathcal{X},\mathcal{X}'}) = (1-s)\cdot a_{f(\mathcal{X}\cap\mathcal{X}')}(\delta_{\mathcal{X}\cap\mathcal{X}'}) + s\cdot a_{f(\mathcal{X}\cup\mathcal{X}')}(\nu_{\mathcal{X}\cup\mathcal{X}'})$.

Proof. If A and B are orthogonal, then

$$a_{f(\mathcal{X})}(s\,A \oplus (1-s)\,B) = a_{f(\mathcal{X})}(s\,A) + a_{f(\mathcal{X})}((1-s)\,B).$$

Therefore, since $\delta_{\mathcal{X}\cap\mathcal{X}'}$ and $\nu_{\mathcal{X}\cup\mathcal{X}'}$ are orthogonal, we have $a_{f(\mathcal{X})}(\mu^s_{\mathcal{X},\mathcal{X}'}) = (1-s)\cdot a_{f(\mathcal{X})}(\delta_{\mathcal{X}\cap\mathcal{X}'}) + s\cdot a_{f(\mathcal{X})}(\nu_{\mathcal{X}\cup\mathcal{X}'})$.

Since $f(\mathcal{X}), f(\mathcal{X}\cap\mathcal{X}') \in \operatorname{aff}(\nu_{\mathcal{X}\cap\mathcal{X}'})$ and $\delta_{\mathcal{X}\cap\mathcal{X}'}$ is orthogonal to $\nu_{\mathcal{X}\cap\mathcal{X}'}$, we have $a_{f(\mathcal{X})}(\delta_{\mathcal{X}\cap\mathcal{X}'}) = a_{f(\mathcal{X}\cap\mathcal{X}')}(\delta_{\mathcal{X}\cap\mathcal{X}'})$. Similarly, $f(\mathcal{X}), f(\mathcal{X}\cup\mathcal{X}') \in \operatorname{aff}(\delta_{\mathcal{X}\cup\mathcal{X}'})$ and $\delta_{\mathcal{X}\cup\mathcal{X}'}$ is orthogonal to $\nu_{\mathcal{X}\cup\mathcal{X}'}$, hence $a_{f(\mathcal{X})}(\nu_{\mathcal{X}\cup\mathcal{X}'}) = a_{f(\mathcal{X}\cup\mathcal{X}')}(\nu_{\mathcal{X}\cup\mathcal{X}'})$. Concluding, $a_{f(\mathcal{X})}(\mu^s_{\mathcal{X},\mathcal{X}'}) = (1-s)\cdot a_{f(\mathcal{X}\cap\mathcal{X}')}(\delta_{\mathcal{X}\cap\mathcal{X}'}) + s\cdot a_{f(\mathcal{X}\cup\mathcal{X}')}(\nu_{\mathcal{X}\cup\mathcal{X}'})$. □

Now that we have the decomposition of the anchor point of a polyhedral cell into the anchor point of Delaunay and Voronoi cells, we show that these anchor points are easily constructed.

Lemma 53. *The anchor point* $a_{f(\mathcal{X})}(\delta_{\mathcal{X}})$ *lies in the interior of* $\delta_{\mathcal{X}}$ *or* $a_{f(\mathcal{X})}(\delta_{\mathcal{X}}) = a_{f(\mathcal{X}')}(\delta_{\mathcal{X}'})$, *where* $\delta_{\mathcal{X}'}$ *is a face of* $\delta_{\mathcal{X}}$

Proof. Assume that $f(\mathcal{X})$ is not contained in $\text{int}(\delta_{\mathcal{X}})$, otherwise the proof is trivial. Since $\delta_{\mathcal{X}}$ is a convex polyhedron, the point closest to $f(\mathcal{X})$ lies on a proper face of $\delta_{\mathcal{X}}$, say $\delta_{\mathcal{X}'}$.

Since $f(\mathcal{X}), f(\mathcal{X}') \in \text{aff}(\nu_{\mathcal{X}'})$ and $\delta_{\mathcal{X}'}$ is orthogonal to $\nu_{\mathcal{X}'}$, we have $a_{f(\mathcal{X})}(\delta_{\mathcal{X}'}) = a_{f(\mathcal{X}')}(\delta_{\mathcal{X}'})$. □

A similar lemma holds for Voronoi cells, for which we omit the proof.

Lemma 54. *The anchor point $a_{f(\mathcal{X})}(\nu_{\mathcal{X}})$ lies in the interior of $\nu_{\mathcal{X}}$ or $a_{f(\mathcal{X})}(\nu_{\mathcal{X}}) = a_{f(\mathcal{X}')}(\nu_{\mathcal{X}'})$, where $\nu_{\mathcal{X}'}$ is a face of $\nu_{\mathcal{X}}$*

Concluding, from the anchor points $a_{f(\mathcal{X})}(\delta_{\mathcal{X}})$ and $a_{f(\mathcal{X})}(\nu_{\mathcal{X}})$, with $\delta_{\mathcal{X}} \in \text{Del}(\mathcal{P})$ we can construct the anchor point of any polyhedral cell. Moreover, Lemma 53 and Lemma 54 give a recursive definition that makes it easy to compute $a_{f(\mathcal{X})}(\delta_{\mathcal{X}})$ and $a_{f(\mathcal{X})}(\nu_{\mathcal{X}})$.

5.3 The meshing algorithm

This section describes the construction of a tetrahedral complex for which the intersection of a cell with the skin surface is either empty or a topological disk. Moreover we show that the mesh extracted from this tetrahedral complex by the marching tetrahedra algorithm [92] is isotopic to the skin surface.

5.3.1 Monotonicity condition

In Section 5.3.2 we give a detailed construction of the tetrahedral complex. For now, we only give the main condition imposed on the tetrahedral complex. First, we require that each tetrahedron is contained in a single mixed cell. Recall that the skin surface restricted to a mixed cell $\mu^s_{\mathcal{X}}$ is a subset of the quadric $I_{\mathcal{X}}^{-1}(0)$, cf. Equation (5.1). Express a point $x = (x_1, x_2, x_3)$ in the local coordinate system of $I_{\mathcal{X}}$.

Condition 55 (Monotonicity). *Let ab be a line segment contained in a mixed cell $\mu^s_{\mathcal{X}}$ of type ℓ, with $I_{\mathcal{X}}(a) \leq I_{\mathcal{X}}(b)$. The segment ab satisfies the monotonicity condition if $x_1^2 + \ldots + x_\ell^2$ is non-increasing and $x_{\ell+1}^2 + \ldots + x_3^2$ is non-decreasing on the segment from a to b.*

In words, a segment ab satisfies the monotonicity condition if the distance to both symmetry sets of the quadric $I_{\mathcal{X}}$ is monotone and the distance to one symmetry set of $I_{\mathcal{X}}$ does not increase if the distance to the other symmetry set increases. For spheres ($\ell = 0, 3$) one symmetry set is empty and the monotonicity condition is satisfied if the distance to the center of the sphere is monotone. For hyperboloids ($\ell = 1, 2$) a segment satisfies the monotonicity condition if the distances to the symmetry axis and the symmetry plane are monotone and the distance to one symmetry set does not increase if the distance to the other symmetry set increases. From Equation (5.1) we conclude:

Observation 56. *If a line segment* $\mathtt{ab}$ *satisfies the monotonicity condition, then* $I_{\mathcal{X}}$ *is monotonically increasing on* $\mathtt{ab}$.

We construct the tetrahedral complex in such a way that all edges satisfy the monotonicity condition. In fact, if all edges satisfy the monotonicity condition, then a generalized monotonicity condition holds for all cells.

Lemma 57. *Let* $\mu^s_{\mathcal{X}}$ *be a mixed cell of type* ℓ *and let* $v_1, \ldots, v_n$ *be the vertices of a cell of the tetrahedral complex in* $\mu^s_{\mathcal{X}}$, *with* $I_{\mathcal{X}}(v_i) \leq I_{\mathcal{X}}(v_j)$ *if* $i < j$.

If the monotonicity condition holds for all edges then, each segment $\mathtt{ab}$, *with* $\mathtt{a} \in \mathrm{conv}(v_1, \ldots, v_k)$ *and* $\mathtt{b} \in \mathrm{conv}(v_{k+1}, \ldots, v_n)$, *for* $k \in \{1, \ldots, n\}$, *satisfies the monotonicity condition.*

In the proof we need a small lemma.

Lemma 58. *Let* $v_1v_2v_3$ *be a triangle in* $\mathbb{R}^2$, *such that the distance to the origin is monotonically increasing (decreasing) along both* v_1v_3 *and* v_2v_3. *Then the distance to the origin is monotonically increasing (decreasing) on the segment* xv_3, *with* $x \in v_1v_2$.

Proof. Let $x = (1-t)v_1 + tv_2$, with $t \in (0,1)$ and let $d(\gamma) = \| (1-\gamma)x + \gamma v_3 \|^2$ be the squared distance to the origin on the line segment xv_3. The distance $d(\gamma)$ is monotone if $d'(\gamma) \geq 0$ or $d'(\gamma) \leq 0$ for $\gamma \in [0,1]$. Since $d'(\gamma) = 2\langle x, v_3 - x\rangle + 2\gamma\langle v_3 - x, v_3 - x\rangle$, $d(\gamma)$ is monotone if $\langle x, v_3 - x\rangle \geq 0$ or $\langle v_3, v_3 - x\rangle \leq 0$.

Assume that the distance to the origin increases monotonically on both line segments v_1v_3 and v_2v_3, hence $\langle v_1, v_3 - v_1\rangle \geq 0$ and $\langle v_2, v_3 - v_2\rangle \geq 0$. We have

$$\begin{aligned} \langle x, v_3 - x\rangle &= (1-t)\langle v_1, v_3 - v_1\rangle \\ &\quad + t\langle v_2, v_3 - v_2\rangle \\ &\quad + t(1-t)\| v_2 - v_1 \|^2 \\ &\geq 0. \end{aligned}$$

Conversely assume that the distance to the origin decreases monotonically on both line segments v_1v_3 and v_2v_3, hence $\langle v_3, v_3 - v_1\rangle \leq 0$ and $\langle v_3, v_3 - v_2\rangle \leq 0$. Then we have $\langle v_3, v_3 - x\rangle = (1-t)\langle v_3, v_3 - v_1\rangle + t\langle v_3, v_3 - v_2\rangle \leq 0$. □

Proof of Lemma 57. We repeatedly move vertices along edges of the cell of the tetrahedral complex while maintaining the monotonicity condition. After the displacement of the vertices, the line segment $\mathtt{ab}$ lies on one of the edges.

Since $I_{\mathcal{X}}(v_i) \leq I_{\mathcal{X}}(v_j)$ for $i < j$ and by Equation (5.1), the distance to the symmetry set spanned by the first ℓ coordinate axis is decreasing and the distance to the other symmetry set is increasing. Assume for now that if we move the vertices v_i and v_{i+1} over the edge v_iv_{i+1}, that all edges in the new cell satisfy the monotonicity condition.

We now move v_1 to $\mathtt{a}$ and v_n to $\mathtt{b}$. If $k = 1$, then $\mathtt{a}$ is v_1. Otherwise we move v_{k-1} to v_k until $\mathtt{a}$ lies on the face $\mathrm{conv}(v_1, \ldots, v_{k-1})$ and repeat this step for $k-1$. Similarly we move v_{k+1} to v_{k+2} until $\mathtt{b}$ lies on the face $\mathrm{conv}(v_{k+1}, \ldots, v_n)$ and

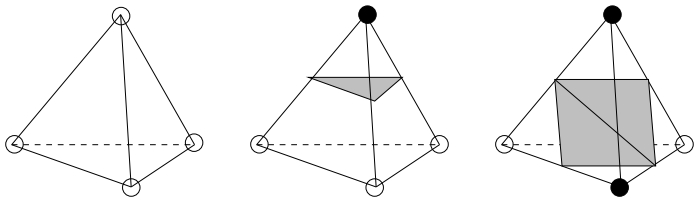

Figure 5.5: The three different configurations of a tetrahedron. White and black vertices lie on different sides of the skin surface.

repeat the step for $k+1$ until $k = n-1$. Then ab is an edge of the new tetrahedron. Hence ab satisfies the monotonicity condition.

It remains to show that the new edges also satisfy the monotonicity condition. Therefore consider three vertices v_i, v_{i+1} and v_j. If $j < i$ ($j > i+1$) then the distance to the first symmetry set along the line segment v_jv_i and v_jv_{i+1} is decreasing (increasing) and to the second symmetry set it is increasing (decreasing). We show that the distances remain monotonically increasing or decreasing along a line segment xv_j for $x \in v_iv_{i+1}$. We distinguish three cases. First, assume that the symmetry set is a point. We project the symmetry point on the plane $v_jv_iv_{i+1}$. From the previous lemma it follows that the distance on xv_j to the projection of the symmetry point is monotone, and therefore also the distance to the symmetry point. Next, assume that the symmetry set is a line, then we project the triangle $v_jv_iv_{i+1}$ on a plane orthogonal to the symmetry line. By applying the previous lemma, it follows that the distance to the projection of the line on the plane along the line segment xv_j is monotone. Finally, if the symmetry set is a plane we do not need the previous lemma. In this case the distance to the plane at x is smaller (greater) than the distance at v_j. Hence the distance to the symmetry plane on xv_j is monotone.

To conclude, if the distance to a symmetry set is monotonically increasing (decreasing) along both v_jv_i and v_jv_{i+1}, then this distance is also monotonically increasing (decreasing) along the line segment xv_j, with $x \in v_iv_{i+1}$. Hence xv_j satisfies the monotonicity condition. □

Mesh extraction The coarse mesh is extracted from the tetrahedral complex by the marching tetrahedra algorithm [92]. Each edge of the tetrahedral complex intersects the skin surface at most once by Observation 56. We place vertices of the mesh on these intersection points. Then the mesh is constructed by considering the number of vertices of the tetrahedron inside the skin surface as depicted in Figure 5.5. The third configuration remains ambiguous, since the common interior edge of the two triangles can be flipped.

Theorem 59. *A tetrahedral complex for which each edge satisfies the monotonicity condition has two properties:*

1. *each cell intersects the skin surface in a topological disk and*

2. *the mesh extracted from the tetrahedral complex is isotopic to the skin surface.*

Proof. Let V^- and V^+ be the vertices of a k-cell of the tetrahedral complex inside and outside the skin surface, respectively. Consider the set of line segments ab with $a \in \text{conv}(V^-)$, $b \in \text{conv}(V^+)$. The set of line segments is empty if the cell does not intersect the skin surface, i.e. if $V^- = \emptyset$ or $V^+ = \emptyset$. On the other hand, if the cell intersects the skin surface, then the set of line segments spans the cell and the line segments may intersect but only at their endpoints. On faces of the cell, the line segments are defined consistently because there the construction is based only on the labels of vertices of the face.

By Lemma 57, each segment satisfies the monotonicity condition. Hence $I_\mathcal{X}$ is monotone on ab. Moreover, a lies inside and b outside the skin surface. Therefore ab intersects the skin surface in a single point. Since the segments span the tetrahedron, the skin surface within the cell is a topological disk.

By construction, each segment also intersects the coarse mesh transversally in exactly one point. We construct the isotopy by constructing an isotopy between the mesh and the skin surface within each tetrahedron and by showing that the isotopies defined by two tetrahedra are identical on a common face.

Neither the skin surface nor the mesh intersects a tetrahedron with only inside (outside) vertices, and we define the isotopy by the identity function. If the skin surface intersects a tetrahedron, then the segments span the tetrahedron and each point in the tetrahedron and on the skin surface lies on a unique segment. The same holds for a point in the tetrahedron and on the mesh. We construct the isotopy by linearly moving each point on the skin surface along the segment to the mesh. By construction of the segments, this deformation is an isotopy.

To show that these local isotopies can be combined to form an isotopy between the skin surface and the coarse mesh it remains to show that the transition between two local isotopies is continuous. This follows from the construction of the segments on a common face, which depends only on the label of its vertices. □

We call the segments in the proof above *transversal segments* because each segment intersects both the skin surface and the coarse mesh transversally in a single point.

5.3.2 The tetrahedral complex

Up to now we assumed that it is possible to construct a tetrahedral complex in such a way that all edges satisfy the monotonicity condition. In this section we construct this tetrahedral complex. We triangulate polyhedral cells in order of increasing dimension.

All vertices of the tetrahedral complex are anchor points of polyhedral cells. In case an anchor point lies on the boundary of its polyhedral cell, it coincides with another anchor point and the simplicial complex is degenerate. Therefore, during the construction of the tetrahedral complex we test whether the anchor point lies in the interior of the polyhedral cell, and collapse the vertex otherwise. For simplicity, in the remainder of this section we assume that the anchor point lies in the interior of the mixed cell.

Subdividing polyhedral cells of positive co-dimension On each vertex of the polyhedral complex we place a vertex of the tetrahedral complex. Note that these vertices are the anchor point of 0-cells of the polyhedral complex.

Next, consider an edge $\mu^s_{\mathcal{X},\mathcal{X}'}$ of the polyhedral complex. By Lemma 50, if $I_{\mathcal{X}}$ has a critical point on the interior of the edge, this critical point is the anchor point of the edge. Therefore we split the edge in the anchor point $a_{f(\mathcal{X})}(\mu^s_{\mathcal{X},\mathcal{X}'})$ and construct two edges from the anchor point to the vertices. By Observation 43 a polyhedral edge is parallel to one symmetry set and is split in the point closest to the projection of the other symmetry set, hence both edges satisfy the monotonicity condition.

We distinguish two types of facets: *circular* facets are facets for which the contour lines of $I_{\mathcal{X}}$ restricted to the facet are circles. The other facets are called *hyperbolic* because the contour lines are hyperbolas on the facet. Since the skin surface is tangent continuous, $I_{\mathcal{X}}|\mu^s_{\mathcal{X},\mathcal{X}'} = I_{\mathcal{X}'}|\mu^s_{\mathcal{X},\mathcal{X}'}$ and the facet inherits the same type from both mixed cells $\mu^s_{\mathcal{X}}$ and $\mu^s_{\mathcal{X}'}$. All facets of mixed cells of type 0 and 3 are spherical. The facets of a mixed cell of type 1 or 2 are spherical if they touch a mixed cell of type 0 or 3, and hyperbolic if they touch a mixed cell of type 1 or 2.

We triangulate circular and hyperbolic facets differently. Circular facets are triangulated by adding an edge from the anchor point of the facet to each anchor point on the boundary of the facet, i.e., either the anchor point of an edge or a vertex. See Figure 5.4(a). Since the anchor point of the facet is the point closest to the center of the sphere, the distance to the focus increases monotonically on each edge and each edge satisfies the monotonicity condition.

Hyperbolic facets are rectangles with edges parallel or perpendicular to the symmetry axis of the corresponding hyperboloid. The anchor point of an edge is the point closest to the focus, hence it is the point on the edge closest to the symmetry axis the edge is orthogonal to. Similarly, the anchor point of the facet is the point on the facet closest to the focus. Thus the edges from the anchor point of the facet to the anchor point of an edge are parallel to one axis and the distance to the other axis increases monotonically. Further, we add edges from the anchor point of an edge to the anchor point of an orthogonal edge. On these edges the distance to one symmetry axis increases whereas the distance to the other symmetry axis decreases. This triangulation is depicted in Figure 5.4(b).

Subdividing polyhedral cells of type 0 and 3 The mixed cells of this type contain a spherical patch of the skin surface. Similar to spherical facets, we have to triangulate polyhedral cells of type 0 and 3 in such a way that the distance to the focus is monotone on each edge. The anchor point of the mixed cell is the point in the mixed cell closest to the focus. Hence the distance to the focus on each line segment from the anchor point of the mixed cell to any other point in the mixed cell is monotone, and therefore satisfies the monotonicity condition. We have already constructed the triangulation of the boundary of the mixed cell and triangulate the entire cell by adding edges from the anchor point of the cell to each vertex on the boundary. The tetrahedra are formed by taking the join of a triangle on the the triangulated boundary of the mixed cell and the anchor point of the mixed cell.

Subdividing polyhedral cells of type 1 and 2 The triangulation of mixed cells of type 1 and 2 is slightly more subtle. The mixed cell contains a hyperboloid patch of the skin surface and the mixed cell is a prism with its base parallel to the symmetry plane of the hyperboloid. For an edge to satisfy the monotonicity condition, the distance to both the symmetry plane and the symmetry axis has to be monotone and the distance to one symmetry set may not increase, if the distance to the other symmetry set increases.

We already triangulated the facets of the mixed cell. The hyperbolic facets of the prism are the facets that are parallel to the symmetry axis. We split the prism in the plane V through the anchor point of the mixed cell parallel to the symmetry plane. This plane also contains the anchor points of the faces and edges of the mixed cells that are parallel to the symmetry axis. Hence each facet parallel to the symmetry axis is already split in V. The new facet is spherical and we triangulate it accordingly.

Consider one split mixed cell. The base of the prism furthest away from the symmetry plane contains the points furthest away from the symmetry plane. Hence its anchor point is the point with maximal distance to the symmetry plane and minimal distance to the symmetry axis. Therefore, all line segments in the split mixed cell with this anchor point as a vertex satisfy the monotonicity condition. The boundary of the prism is already triangulated and we triangulate the split mixed cell by adding edges from the anchor point of the base to all vertices on the boundary. The tetrahedra are the join of a triangle on the triangulated boundary and the anchor point of the base.

Union of balls For a shrink factor one, the skin surface of a set of balls is the union of these balls. In this case, the mixed complex is the Voronoi diagram. This means that only mixed cells of type 0 are three-dimensional cells. This greatly simplifies the set of tetrahedra.

It is also desirable to retain edges of the mesh on the intersection of two balls. The subdivision algorithm ensures this by definition. Chew's algorithm can also be extended to allow constrained edges, see [34].

5.3.3 Complexity analysis

In many real world applications the size of the Delaunay triangulation is linear in the number of input balls, see [36, 10]. However, the worst case complexity of the Delaunay triangulation is quadratic in the number of input balls, see [22]. We show that the size of the coarse mesh is linear in the size of the Delaunay triangulation and that this is worst case optimal.

Lemma 60. *The size of the coarse mesh is linear in the size of the Delaunay triangulation.*

Proof. Because of the duality relationship, the size of the Voronoi complex is equal to the size of the Delaunay triangulation. Each mixed 0-cell (corresponding to a Delaunay vertex, is a shrunk Voronoi 3-cell, and its size is therefore equal to the

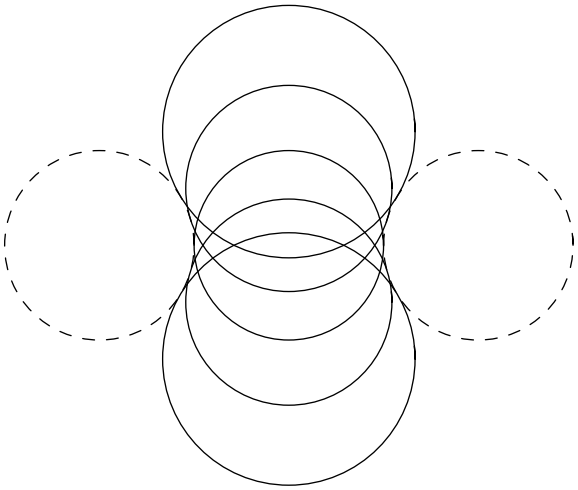

Figure 5.6: The cross section in the xz-plane of the skin surface with a quadatic number of holes. The centers of the dashed circles lie on the unit circle in the xy-plane and the centers of the solid circles on the z-axis.

Voronoi cell. Similarly, the complexity of a mixed 1-cell, 2-cell and 3-cell are linear in the complexity of the Voronoi facet, Delaunay facet and Delaunay tetrahedron.

We split each edge of the mixed complex at most in two parts. The triangulation of a mixed facet contains at most one triangle per split edge. Finally, the triangulation of a mixed cell contains at most one tetrahedron per triangle on the mixed facets. Hence, the size of the tetrahedral complex is linear in the mixed complex.

Within each tetrahedron we construct at most two triangles. Thus, the mesh is linear in the size of the Delaunay triangulation. □

To show that this is worst case optimal, we construct a skin surface with $\Omega(n^2)$ holes from a set of n balls. Any mesh with $\Omega(n^2)$ holes has complexity $\Omega(n^2)$, thus giving the lower bound. The construction, depicted in Figure 5.6, is as follows: the first $n/2$ balls are centered on the unit circle in the xy-plane and have radius 0.5. The other $n/2$ balls are centered on the z-axis. Their radius is such that they touch the first $n/2$ balls.

Each two subsequent spheres on the z-axis form a tunnel with each sphere centered on the unit circle. There are $n/2 - 1$ such pairs and $n/2$ spheres centered on the unit circle, hence there are $\Omega(n^2)$ tunnels. The skin surface also has $O(n^2)$ holes because it is homeomorphic to the union of the balls.

5.3.4 Mesh enhancement

The topologically correct mesh obtained with the marching tetrahedra algorithm is rather coarse and may contain long and skinny triangles. Therefore, we develop a method to enhance the mesh while maintaining the isotopy. The changes to the mesh we allow are local and do not change the topology of the mesh.

Before we change the mesh, we first test whether the isotopy with the skin surface is maintained. Therefore we use the transversal segments as described in the proof of Theorem 59. In fact, we first test whether each transversal line segment intersects the new mesh exactly once. We conclude this section with two examples of mesh refinement algorithms.

Changing the mesh To test whether the isotopy is maintained under a change of the mesh we would have to test whether each transversal line segment intersects the mesh once. We rephrase this in such a way that it is easier to verify.

Let c be a 3-cell in the tetrahedral complex, t a triangle of the mesh intersecting c and V^- and V^+ the vertices of c inside and outside the skin surface, respectively.

Lemma 61. *If for all* t *and* c, V^- *and* V^+ *are separated by the plane through* t *and* V^- *lies in the direction of the inner part of the mesh, then each transversal line segment within* c *intersects the mesh once.*

Proof. Consider a line segment $\mathsf{p}^-\mathsf{p}^+$, with $\mathsf{p}\pm \in \text{conv}(V^\pm)$. Since p^- and p^+ lie on opposite sides of the mesh, $\mathsf{p}^-\mathsf{p}^+$ intersects the mesh at least once. Assume that it intersects $\mathsf{p}^-\mathsf{p}^+$ more than once, then on the second intersection point from p^-, the segment moves from outside the mesh to the inside. Hence. the inner product with the normal is negative. □

We now have an efficient way of testing whether the isotopy of the skin surface and the mesh is maintained. If the test fails, then the mesh is too coarse and we refine the mesh. We show that the refinement succeeds for small triangles.

Lemma 62. *A triangle* t *of the mesh contained in a single tetrahedron* c *can be subdivided in any point* $\mathsf{x} \in \mathsf{t}$ *by moving* x *along the transversal segments to the skin surface.*

Proof. Since x moves along the line segments within c, the line segments through f and the subdivided faces are the same. Thus the new mesh can be obtained from the old mesh by interpolation along the line segments. □

We can also flip an edge of the triangulation if the two adjacent triangles and the new triangles are intersected by the same transversal segments. This condition is similar to the condition in a two-dimensional mesh that an edge can be flipped if the union of the two triangles is convex.

To summarize, we have an efficient test to check whether the isotopy is maintained. If a change of the mesh would result in a violation of the isotopy test, then we can always subdivide the face into faces that are contained within a single tetrahedron. Each of these faces satisfies the isotopy test.

Sqrt-3 method We implemented the sqrt-3 subdivision method [69] on the coarse mesh. The sqrt-3 subdivision method splits each triangle into 9 sub-triangles and then moves the newly created vertices towards the skin surface along the transversal segments.

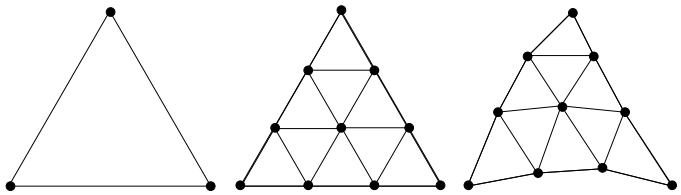

Figure 5.7: The Sqrt-3 subdivision method applied twice. Left: the original triangle, in the middle the subdivided triangle. On the right are the vertices placed on the skin surface.

By Lemma 62, the subdivision algorithm maintains the isotopy. Hence, it is not necessary to test isotopy, which make the algorithm very fast. On the other hand, the subdivision algorithm does not improve the quality of the triangles. Therefore this method is not suitable for constructing a mesh for numerical simulations.

Chew's algorithm We also implemented Chew's algorithm [34] to improve the quality of the triangles of the coarse mesh and obtain a mesh suitable for numerical simulations. After the algorithm terminates, each triangle has angles between 30 and 120 degrees and has a user defined maximal size. The only constraint on the size-criterion is that there exists a $\delta > 0$ such that any well-shaped triangle that fits within a circle of radius δ is well-sized. We chose the size of a triangles inversely proportional to the maximal curvature which is nonzero on skin surfaces.

During the refinement, we test the isotopy before an inserting a new point and before flipping an edge.

5.4 Implementation

We implemented the algorithm described above in C++ using the Computational Geometry Algorithm Library (CGAL) [37]. This algorithm will be available as a CGAL extension package.

CGAL is a library written in C++ that uses generic programming to attain its full flexibility. It can be subdivided in three parts: the kernel, datastructures and algorithms and supporting classes.

The kernel provides geometric primitive objects like points, line segments, triangles etc. The kernel also defines predicates and constructions on these objects. Predicates are functions that test a property (return a boolean). For example the "incirle test" used for the construction of the Delaunay triangulation is a predicate. Constructions are functions that return a geometric object, for example the intersection point of a line and a plane. Several kernels are provided in CGAL: (1) inexact predicates and inexact constructions, (2) exact predicates and inexact constructions and (3) exact predicates and exact constructions.

The second part of CGAL contains a collection of basic geometric data structures and algorithms, which are parameterized by traits classes that define the interface

between the data structure or algorithm and the primitives they use. Sometimes, the datastructures and algorithms are also parameterized by observer classes. Observer classes define functions that are called before or after an event occurs, they "observe" the algorithm.

The support library contains non-geometric support facilities, such as circulators, I/O support and interfaces between CGAL and various visualization tools, like geomview and QT.

5.4.1 Outline of the implementation

For the implementation of the algorithm the kernel is extended with a quadratic surface class to store the part of the skin surface contained inside a mixed cell.

We did not introduce new datastructures, but customized the datastructures available in CGAL to suit our needs. The weighted Delaunay triangulation is implemented in the `Regular_triangulation_3` class and does not need adaptations. The triangulated mixed complex `Triangulated_mixed_complex_3` is derived from the triangulation class `Triangulation_3` and has cells that additionally store a pointer to the quadratic surface defined inside the cells. The triangulated mesh is stored as a `Polyhedron_3`.

We implemented the function `triangulate_mixed_complex_3` that constructs the triangulated mixed complex from the Delaunay triangulation and the function `marching_tetrahedra_3` that extracts the triangulated mesh from the triangulated mixed complex. Further, we implemented the subdivision algorithm in the function `skin_surface_sqrt3` and a prototype of Chew's algorithm. An outline of the implementation is depicted in Figure 5.8.

No functionality was added to the support library.

In the remainder of this section we discuss the implementation in more detail. The kernels described below are the default kernels, everything can be customized using techniques like traits classes and templates.

Delaunay triangulation.

We use the `Regular_triangulation_3` to compute the weighted Delaunay triangulation of the input balls. Since no geometric objects are constructed during the construction of the Delaunay triangulation, any kernel with exact predicates produces the correct Delaunay triangulation.

Triangulated mixed complex.

The triangulated mixed complex is a `Triangulation_3`. Its kernel should also be able to do exact constructions, since the vertices of the triangulation are anchor points of simplices of the Delaunay triangulation, which are constructed. To speed up the computation, we use a filtering technique [24] that is present in CGAL. This filter first computes small boxes containing the vertices and uses those in further computations. If the filtered computations fail to give a guaranteed result, the computation is redone with an exact number type.

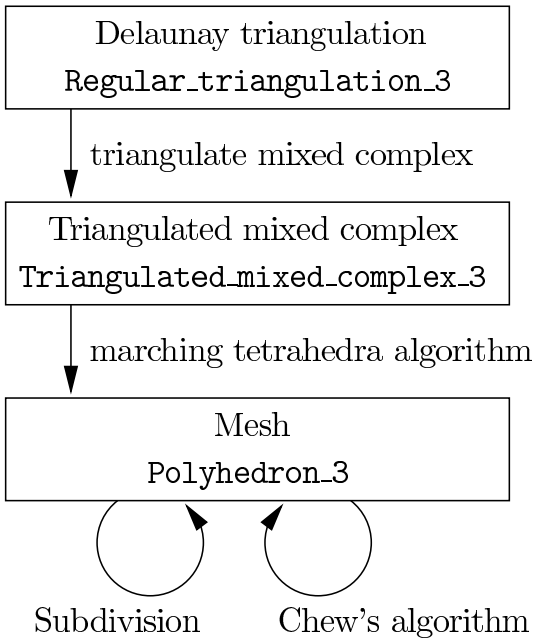

Figure 5.8: Outline of the implementation. The boxes contain the datastructures (with the name of the datastructure above the name of the class) and the labeled arrows denote the functions that construct one datastructure from the other. The labels describe the functions.

By definition, every tetrahedron in the triangulation is contained in a single mixed cell. We store a pointer to the quadratic surface corresponding to the mixed cell with each cell (tetrahedron). In this way we are able to move a point contained in the cell towards the skin surface. This is done by first constructing the transversal segment containing the point form the cell and then intersecting the transversal segment with the quadratic surface.

Mesh

For the mesh, we use the `Polyhedron_3` class which is able to store any orientable 2-manifold. The kernel has to be able to perform exact constructions since its vertices are constructed. Further the kernel has to support the predicates exactly that are needed in the refinement step.

Construction of the triangulated mixed complex

The construction of the triangulated mixed complex is the least straightforward part of the algorithm. It uses a `Triangulation_incremental_builder_3` that is developed for this algorithm, but might be of general interest. This builder first gets the vertices of the triangulation and then the cells, which are defined by references to four vertices. From these vertices and cells it constructs the triangulation.

The function `triangulate_mixed_complex_3` computes the vertices and cells of the triangulation and inserts them in the `Triangulation_incremental_builder_3`.

The construction of vertices is complicated due to the occurrence of many degenerate cases. For example, let $\nu_{\mathcal{X}}$ be a Voronoi edge and let $\nu_{\mathcal{X}'}, \nu_{\mathcal{X}''}$ be its two vertices. If the Voronoi edge degenerates into a point, then the anchor points $\mathfrak{a}_{\mathfrak{f}(\mathcal{X})}(\nu_{\mathcal{X}})$, $\mathfrak{a}_{\mathfrak{f}(\mathcal{X})}(\nu_{\mathcal{X}'})$ and $\mathfrak{a}_{\mathfrak{f}(\mathcal{X})}(\nu_{\mathcal{X}''})$ are equal. Similar cases occur if a Voronoi facet degenerates into a line segment or a Voronoi cell into a facet. We first compute the anchor points of the triangulated mixed complex and combine anchor points that are equal using a union-find datastructure. We then insert all unique vertices and the tetrahedra that are not degenerate, i.e., have four distinct vertices.

A similar class exists for the triangulation of the Voronoi diagram. This triangulation can be used for meshing the boundary of the union of a set of balls.

Marching tetrahedra algorithm

The marching tetrahedra algorithm `marching_tetrahedra_3` introduced in [92] takes four arguments: the triangulation, the polyhedron, a traits class and an observer class. The algorithm is performed on the triangulation and stored in the polyhedron. The `Polyhedron_incremental_builder_3` is used for constructing the polyhedron.

The traits class defines a single predicate and a construction. First it is able to test whether a vertex of the triangulation lies inside or outside the surface and it is able to return an intersection point of the surface with an edge of the triangulation whose vertices lie on opposite sides of the surface. For skin surfaces the intersection point is unique as is shown in Theorem 59.

The observer class implements two functions that are called after the construction of a vertex and a facet of the polyhedron. After insertion of a vertex in the polyhedron a function is called with the vertex of the polyhedron and the corresponding edge of the triangulation. Similarly, after insertion of a facet in the polyhedron a function is called with the facet of the polyhedron and the corresponding cell of the triangulation. Using these two functions it is possible to construct a reference from simplices of the mesh to the corresponding simplices of the triangulation.

Both the traits class and the observer class make this algorithm highly flexible.

Refinement algorithms

We also implemented both algorithms described in the Section 5.3.4. The implementation of the $\sqrt{3}$-subdivision method is straightforward. For the position of the new vertices, we use point location in the triangulated mixed complex and then project the new vertex along the transversal segment towards to the skin surface.

Initially, the performance of the subdivision algorithm was disappointing since the filtered point location step failed often leading to slow running times. This failure can be explained since new vertices of subdivided polyhedron lie on edges of the coarse mesh, which in turn lie in facets of the triangulation. Filtered arithmetic often fails in degenerate situations like this. To give the algorithm a significant speedup, we also store with each facet of the polyhedron a pointer to the cell of the triangulation entirely containing the facet. In this way the point location step

Molecule	Our algorithm			Dynamic	Marching
	Coarse	Sqrt-3	Chew		
pdb7tmn	0:00:01	0:00:02	0:00:05	0:10:00	0:00:05
DNA	0:00:14	0:00:29	0:00:55	0:35:12	0:00:51
Gramacidin A	0:00:08	0:00:31	0:01:13	1:35:23	0:03:22

Table 5.1: Performance comparison

is not necessary. Storing this pointer is done with an adapted observer class of the marching tetrahedra algorithm.

The implementation of Chew's algorithm is the only part that is still in a developmental stage. Before we apply Chew's algorithm we perform a preprocessing step in which we remove small edges. This reduces the size of the final mesh considerably.

5.5 Examples and experiments

We compare our algorithm to the algorithms described in [27] and [29]. There is a comparison of the two algorithms in [29]. These tests are run on a Pentium 4 running at 2.54GHZ. To test our algorithm we used an AMD Athlon 1800+ which is actually a little slower. We tested our algorithm on various molecules, computing only the coarse mesh, computing the coarse mesh and after that one $\sqrt{3}$-subdivision step and the coarse mesh which was subsequently improved using Chew's algorithm. For timings see Table 5.1.

Note that both our algorithm and the marching algorithm [29] are significantly faster than the dynamic skin algorithm [27]. However, [29] does not come with topological guarantees.

Figure 5.10 shows the molecule pdb7tmn. In Figure 5.10(d) we enlarged a part of the coarse mesh and applied the Sqrt-3 method in Figure 5.10(e). Note that the triangles remain skinny. Figure 5.10(f) shows the result of applying Chew's algorithm directly to the coarse mesh. Because of small edges in the coarse mesh, there are also small edges near parts with low curvature. If we remove small edges, as in Figure 5.10(g), before we apply Chew's algorithm, we obtain Figure 5.10(h).

5.6 Extension to envelope surfaces

We think that it is possible to adapt the algorithm presented in this chapter to mesh envelope surfaces, introduced in Chapter 4.

The main change lies in the construction of the triangulation of the generalized mixed complex defined in Section 4.4, Proposition 34. Similar to mixed complex of the skin surface, this polyhedral complex decomposes the envelope surface into pieces of quadrics. The function defined within a cell (of any dimension) of the polyhedral complex is a quadric. Any (non-degenerate) quadric has a unique critical point. We

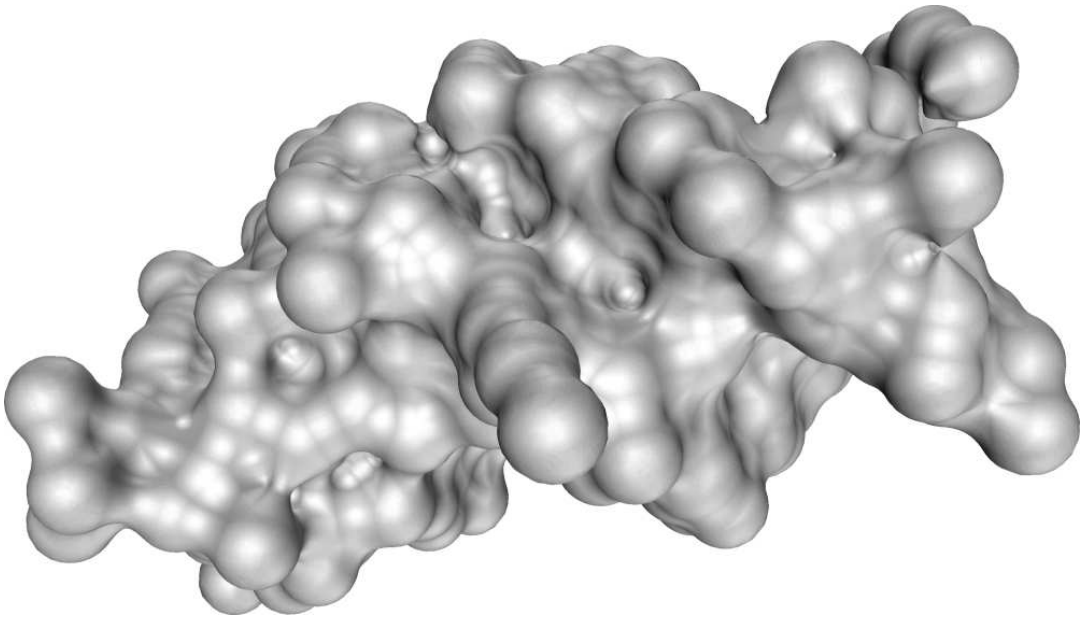

(a) DNA

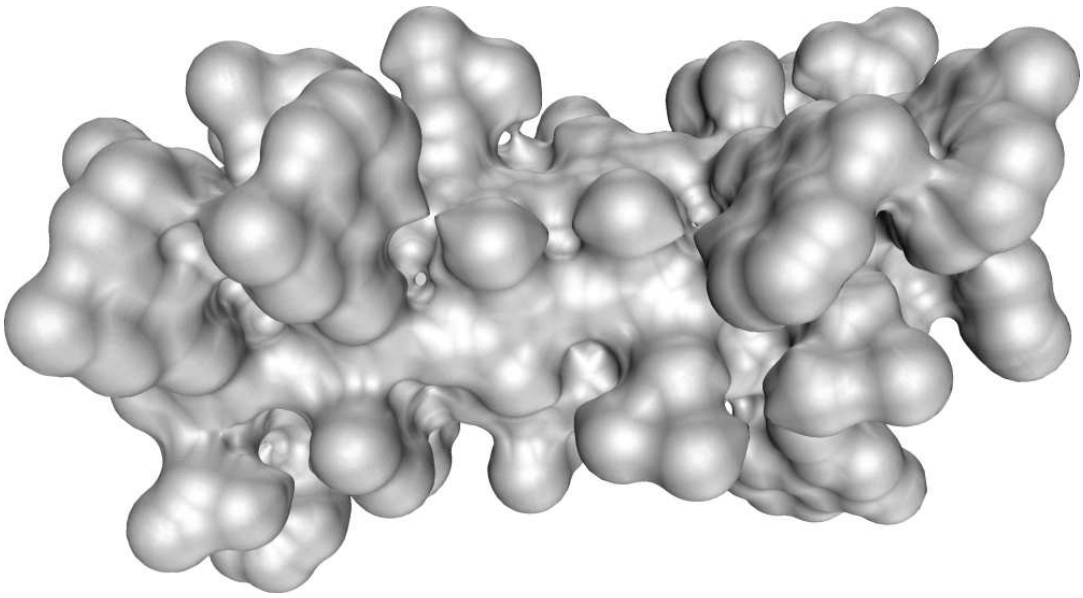

(b) Gramacidin A

Figure 5.9: Two larger molecules.

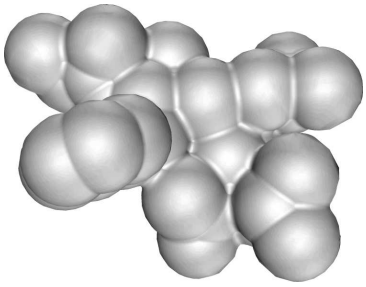

(a) Shrink: .85

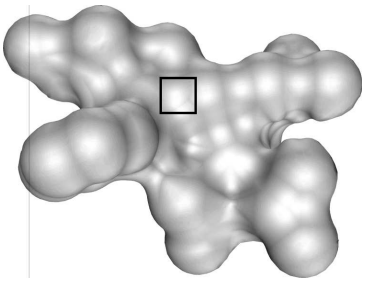

(b) Shrink: .5

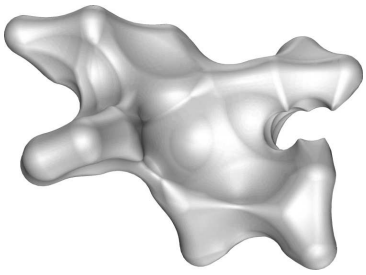

(c) Shrink: .15

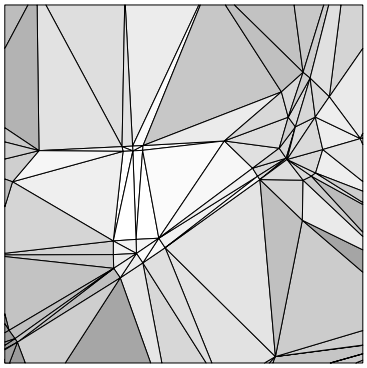

(d) Coarse mesh

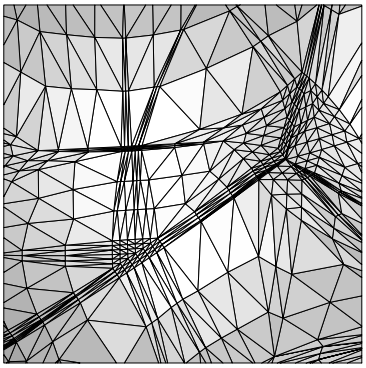

(e) Sqrt3-method

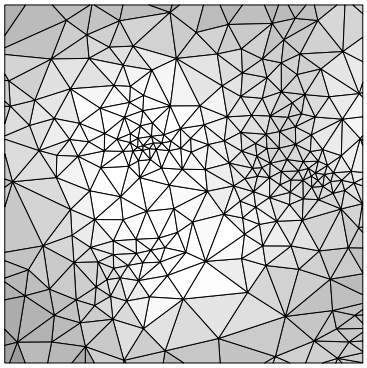

(f) Chew applied to the coarse mesh

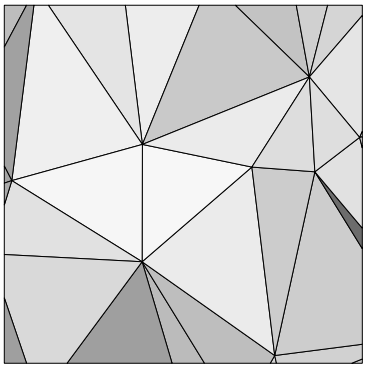

(g) Small edges removed

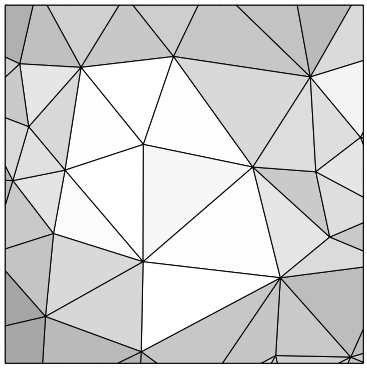

(h) Chew applied to the enhanced mesh

Figure 5.10: The molecule pdb7tmn.

define the vertices of the triangulated polyhedral complex as the anchor points of polyhedral cells with respect to the critical point of the quadric.

Next, we construct a triangulation such that all edges have the monotonicity condition. Therefore, we write the quadric in normal form using the eigenvectors as its basis and we split the polyhedral cell along the planes through the critical points spanned by two eigenvectors. For each non-empty clipped polyhedral cell, we first triangulate the facets of lower dimensional simplices and then construct interior edges. For the triangulation of facets we repeat the algorithm. The triangulation of the interior of the polyhedral cell is done such that on each edge the distance to a basisvector is non-decreasing (non-increasing) along the edge if the corresponding eigenvalue is positive (negative). This implies that all edges satisfy the monotonicity condition and that it is possible to construct the transversal segments. Using these segments we are able to prove Theorem 59 for the general case of envelope surfaces.

5.7 Conclusion and future work

We present an algorithm that constructs a mesh that is isotopic to the skin surface and discuss two methods to refine this mesh.

The algorithm we present is static in the sense that it generates a mesh for a fixed set of input balls.

Maintaining the coarse mesh while deforming the input set, is important for animations and deforming molecules. Two deformation schemes seem computationally interesting. From Lemma 52 we know that the anchor point of a mixed cell $\mu^{\mathrm{s}}_{\mathcal{X},\mathcal{X}'}$ only depends on the Delaunay cell $\delta_{\mathcal{X}\cup\mathcal{X}'}$, the Voronoi cell $\nu_{\mathcal{X}\cap\mathcal{X}'}$ and the shrink factor. Adding a constant to all weights does not change the Delaunay and Voronoi diagram and hence does not change the simplicial complex. Hence, the coarse mesh of the skin surface obtained by adding a constant to each weight is another level-set of the tetrahedral complex. Another deformation is obtained by varying the shrink factor. Again, the structure of the mixed complex and simplicial complex remains unchanged, however the positions of the anchor points change. It is sufficient to reposition the anchor points and then update or recompute the coarse mesh. In general, the solid foundation of the simplicial complex makes us believe that it is possible to maintain the simplicial complex if the Delaunay triangulation can be maintained.

We also believe that the algorithm can be generalized to mesh envelope surfaces that are presented in Chapter 4. For this purpose, we have to implement the changes proposed in Section 5.6. These changes do not seem to make the algorithm much more complicated and is interesting to implement the algorithm in the future and see how this algorithm performs for envelope surfaces.

Similar adaptations make the algorithm suitable for meshing Connelly surfaces. These surfaces are also used in molecular biology and are formed by a small probe sphere that carves away the space outside a union of balls. The disadvantage of this type of surfaces is that they may not be tangent continuous.

Chapter 6

Conclusion and open problems

Conclusion In this thesis we develop a class of C^1 surfaces, called envelope surfaces, that are defined as the envelope of a set of spheres. The set of spheres is constructed by assigning a weight (i.e., squared radius) to a convex and compact set of sphere-centers. The envelope surface is manipulated by changing the weight. Moreover, if the weight function interpolates the weights at centers of a given finite set of input balls, then the envelope surface contains the input balls.

Envelope surfaces extend a class of surfaces called skin surfaces. Skin surfaces [48] were initially used for modeling molecular surfaces. In Chapter 3, we adapt the theory of skin surface to make them suitable for surface approximation. Under certain conditions on the input set of balls, we proved several properties of the approximation with respect to the surface that is approximated: the two surfaces are homeomorphic, the Hausdorff distance is arbitrarily small and every maximal ball in the input set is also maximal in the approximation. For visualization and further geometric processing of skin surfaces we developed a meshing algorithm in Chapter 5. The approximating mesh is isotopic to the skin surface.

Open problems The flexibility of envelope surfaces makes us believe that these surfaces are well suited for the design of surfaces. The approximation with skin surfaces results in surfaces that are bumpy, i.e., the patches blending spheres together are always concave. Using envelope surfaces the patches can be both concave and convex. We conjecture that it is possible to construct an approximation scheme in which the normals of the envelope surface converge to the normals of the original surface.

Except for some prototypes used to give a proof of concept, we have not developed a modeler for these surfaces. With such a modeler it is possible to construct and modify an envelope surface interactively and test the applicability of envelope surfaces to their full extent. Several schemes are proposed in Section 4.6 to construct a weight function interactively. Validation of these schemes on realistic data is also possible with a modeler.

Several deformation schemes have been proposed to deform skin surfaces [28, 33, 48]. These algorithms have been developed for deformations of molecules. We think

that these algorithms can also be used to deform envelope surfaces, which is important in the setting of animations. It would be interesting to see how these schemes perform on envelope surfaces. It might be necessary to develop new deformation schemes to obtain better interpolation results.

Since the input of the surface approximation algorithms is a set of spheres contained inside the surface, the construction of these sets is important for the applicability of the approximation algorithms. Algorithms exist that construct a set of balls from a dense point set on the surface. To the best of our knowledge, there are no general algorithms that construct a set of maximal balls directly from a smooth surface. These algorithms are of interest to us since they can be used to generate the input balls.

We think that, with minor changes, the meshing algorithm for skin surfaces can be generalized to mesh envelope surfaces. Also algorithms maintaining a skin surface or an envelope surface of a deforming set of spheres are of great interest if animating with skin surfaces or envelope surfaces.

Appendix A

Convex functions and differentiability

We review some basic concepts and properties from the theory of differentiable functions (See, e.g., [9]) and from Convex Analysis (See, e.g., [16] and [88]).

A.1 Maxima of parameterized families

In this section X is a topological space, D is a *compact* topological space, and $\Phi : X \times D \to \mathbb{R}$ is a continuous function. The function Φ may also be considered as a *family of functions* $\Phi_p : X \to \mathbb{R}$, *parameterized* by $p \in D$, where $\Phi_p(x) = \Phi(x, p)$. We will also use the shorthand notation $\Phi(\cdot, p)$ for Φ_p. Furthermore, we consider the maximum-function $\varphi : X \to \mathbb{R}$, defined by

$$\varphi(x) = \max_{p \in D} \Phi(x, p).$$

Lemma 63. *The function φ is continuous.*

Proof. Let $x \in X$, and let $\varepsilon > 0$. We shall prove that there is a neighborhood X_0 of x in X such that $|\varphi(\xi) - \varphi(x)| \leq \varepsilon$ for $\xi \in X_0$. For $p \in D$ there is a neighborhood X_p of x in X and a neighborhood D_p of p in D such that

$$|\Phi(\xi, \eta) - \varphi(x)| \leq \varepsilon,$$

for $(\xi, \eta) \in X_p \times D_p$. Since D is compact, the cover $\{D_p \mid p \in D\}$ of D contains a finite sub cover. In other words, there are $p_1, \ldots, p_n \in D$ such that $D = D_{p_1} \cup \cdots \cup D_{p_n}$. Let $X_0 = X_{p_1} \cap \cdots \cap X_{p_n}$, then X_0 is a neighborhood of x in X such that

$$|\Phi(\xi, \eta) - \varphi(x)| \leq \varepsilon,$$

for $\xi \in X_0$ and all $\eta \in D$. It follows that $|\varphi(\xi) - \varphi(x)| \leq \varepsilon$ for $\xi \in X_0$. In other words, φ is continuous in x. □

In general, for given $x \in X$ the value of p for which $\Phi(x, p)$ attains its maximum is not unique. However, if it is unique it depends continuously on x.

Lemma 64. *If for each* x *in* X *there is a unique value* $\lambda(x) \in D$ *such that*

$$\varphi(x) = \Phi(x, \lambda(x)),$$

then the function $\lambda : X \to D$ *is continuous.*

Proof. Let $x \in X$, and consider a sequence $\{x_n\}$ of points in X, converging to x. We shall prove that the sequence $\{\lambda(x_n)\}$ has $\lambda(x)$ as its unique limit point, which allows us to conclude that λ is continuous at x.

Let $\eta \in D$ be a limit point of $\{\lambda(x_n)\}$. Such a limit point exists, since D is compact. Passing to subsequences if necessary we may assume that the sequence converges to η. Since Φ is continuous by assumption, and φ is continuous according to Lemma 63, taking limits in

$$\varphi(x_n) = \Phi(x_n, \lambda(x_n))$$

we obtain the identity

$$\varphi(x) = \Phi(x, \eta).$$

Since the point at which $\Phi(x, \cdot)$ attains its maximum is unique, we conclude that $\eta = \lambda(x)$. In other words, $\lambda(x)$ is the unique limit point of the sequence $\{\lambda(x_n)\}$. □

A.2 Gateaux- and Fréchet-differentiability

Consider a function $f : \mathbb{R}^d \to \mathbb{R}$. The one-sided *directional derivative* of f at $x \in \mathbb{R}^d$ in the direction v, $v \in \mathbb{R}^d$, is

$$f'(x; v) = \lim_{h \downarrow 0} \frac{f(x + hv) - f(x)}{h},$$

provided this limit exists.

The function f is called *Gateaux differentiable* at x if the two-sided directional derivative in the direction v

$$\lim_{h \to 0} \frac{f(x + hv) - f(x)}{h}$$

exists for all $v \in \mathbb{R}^d$. In this case $f'(x; v) = -f'(x; -v)$. The function f is called *Fréchet differentiable* at $x \in \mathbb{R}^d$ if there is a linear function $T_x : \mathbb{R}^d \to \mathbb{R}$ such that

$$f(x + v) = f(x) + T_x(v) + \|v\| E(v),$$

for all v is some neighborhood of x, where $E : R^d \to \mathbb{R}$ is a function such that $E(v) \to 0$ if $\|v\| \to 0$. The linear transformation T_x is called the *total derivative* of f at x, and is denoted by $f'(x)$. If f is Fréchet-differentiable at x, then f is Gateaux-differentiable at x in all directions v, and $f'(x; v) = f'(x)(v)$ for all $v \in \mathbb{R}^d$. Conversely, if f is Gateaux-differentiable at x, and all its partial derivatives are continuous at x, then f is Fréchet-differentiable at x.

Note that this last property does not necessarily hold if the partial derivatives are not continuous. To see this, consider the function $f : \mathbb{R}^2 \to \mathbb{R}$, defined by

$$f(x,y) = \begin{cases} \dfrac{xy^2}{x^2+y^4}, \text{ if } x \neq 0, \\ 0, \text{ if } x = 0. \end{cases}$$

Then f is Gateaux-differentiable at $x_0 = (0,0)$. In fact, if $v = (a,b)$ then $f'(x_0;v) = \frac{b^2}{a}$, if $a \neq 0$, and $f'(x_0;v) = 0$, if $a = 0$. On the other hand f is not even continuous at $(0,0)$, since $f(y^2,y) = \frac{1}{2}$, if $y \neq 0$, and $f(0,0) = 0$. In particular, f is not Fréchet-differentiable at $(0,0)$.

A.3 Convex functions

We continue the discussion of Appendix A.1, but now we assume that $X = \mathbb{R}^d$ and D is a compact subset of $\mathbb{R}^n$. First we recall some results and terminology from Convex Analysis.

A set A is convex if for any two points $x, x' \in A$ the line segment xx' lies in A and A is called strictly convex if the open line segment (x,x') lies in the interior of A for any two points $x, x' \in A$.

The set of convex functions on a compact subset of a Euclidean space is closed.

Lemma 65. *If a sequence of convex continuous functions defined on a compact subset* D *of* $\mathbb{R}^d$ *is convergent, then the limit function is convex on* D.

The proof is straightforward. Note that the result is not true if we restrict to strictly convex functions.

The *epigraph* of a function $f : X \to \mathbb{R}$ is the set of points in $X \times \mathbb{R}$ above the graph of f:

$$\text{epi}(f) = \{(x,z) \mid x \in X, z \in \mathbb{R} \text{ and } f(x) \leq z\}.$$

A function is (strictly) convex if and only if its epigraph is a (strictly) convex set.

The following criterion for strict convexity is useful in the context of piecewise smooth functions.

Lemma 66. *Let* f *be a convex continuous function defined on a closed interval* I *of the real line. If* f *is strictly convex on each connected component of the complement of a finite set of points in* I, *then* f *is strictly convex on* I.

Proof. Let S be the finite set of points such that f is strictly convex on each connected component of $I \setminus S$. Consider a point $\xi \in S$, and a neighborhood U of ξ in I that contains no other points of S. It is sufficient to prove that f is strictly convex on U, since a function is strictly convex on I if every point of I has a neighborhood on which f is strictly convex.

Let $x_0, x_1 \in U$ with $x_0 < x_1$, and let $x_t = (1-t)x_0 + tx_1$ for $0 < t < 1$. Since f is convex on $[x_0, x_1]$, we know that

$$f(x_t) \leq (1-t)f(x_0) + tf(x_1),$$

for $0 < t < 1$. If there is a t, with $0 < t < 1$, for which equality holds, then equality holds for all t. In that case f is an affine function on $[x_0, x_1]$. This contradicts the fact that f is strictly convex on (x_0, ξ) and on (ξ, x_1). Hence the inequality is strict for all t with $0 < t < 1$. In other words, f is strictly convex on U. □

Lemma 67. *If f is C^2 on a convex domain D, then f is (strictly) convex iff the Hessian of f is nonnegative definite (positive definite).*

Proof. If the Hessian of f is nonnegative definite then the function is locally strictly convex. A function that is everywhere locally strictly convex, is strictly convex. □

We present a simple situation in which the maximum-function φ, introduced in Section A.1, is convex.

Lemma 68. *If $\Phi(\cdot, p) : \mathbb{R}^d \to \mathbb{R}$ is convex, for all $p \in D$, then the function $\varphi : \mathbb{R}^d \to \mathbb{R}$ is convex.*

Proof. Note that $(x, z) \in \mathrm{epi}(\varphi)$ if and only if $\varphi(x) \leq z$, i.e., if and only if $\Phi(x, p) \leq z$ for all $p \in D$. Therefore:

$$\mathrm{epi}(\varphi) = \cap_{p \in D} \mathrm{epi}(\Phi(\cdot, p)).$$

The right hand side is an intersection of convex sets, so it is a convex set. □

Lemma 69. *A convex function on $\mathbb{R}^d$ has a (one sided) directional derivative at all points in all directions. Furthermore,*

$$-f'(x; -v) \leq f'(x; v),$$

for all $x, v \in \mathbb{R}^d$.

Proof. Consider a monotonically decreasing sequence h_n of positive numbers tending to 0. Since f is convex, the sequence

$$\frac{f(x + h_n v) - f(x)}{h_n}$$

is monotonically decreasing, so it has a limit. In other words, $f'(x; v)$ exists for all $x, v \in \mathbb{R}^d$. Now let h_- and h_+ be arbitrary positive numbers. Convexity of f implies that

$$\frac{f(x) - f(x - h_- v)}{h_-} \leq \frac{f(x + h_+ v) - f(x)}{h_+}.$$

Taking limits for $h_- \downarrow 0$ and $h_+ \downarrow 0$ yields $-f'(x; -v) \leq f'(x; v)$.

□

Bibliography

[1] M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and C. T. Silva. Point set surfaces. *IEEE Visualization 2001*, pages 21–28, October 2001. ISBN 0-7803-7200-x. [11]

[2] H. Alt and O. Schwarzkopf. The voronoi diagram of curved objects. In *Proceedings of the eleventh Annual ACM Symposium on Computational Geometry*, pages 89–97, New York, NY, USA, 1995. ACM Press. [18]

[3] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete and Computational Geometry*, 22(4):481–504, 1999. [47, 48]

[4] N. Amenta, M. Bern, and D. Eppstein. The crust and the β-skeleton: Combinatorial curve reconstruction. *Graphical Models and Image Processing*, 60:125–135, 1998. [10, 48]

[5] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proceedings of the sixteenth Annual ACM Symposium on Computational Geometry*, pages 213–222, 2000. [10, 11]

[6] N. Amenta, S. Choi, and R.K. Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry: Theory and Applications*, 19:127–153, 2001. [2, 10, 24, 25, 34, 36, 47, 48, 49]

[7] N. Amenta and R.K. Kolluri. Accurate and efficient unions of balls. In *Proceedings of the sixteenth Annual ACM Symposium on Computational Geometry*, pages 119–128, 2000. [2, 25]

[8] F.M. Amirouche. *Principles of Computer Aided Design and Manufacturing*. Prentice Hall, 2nd edition, 2004. [1]

[9] T.M. Apostol. *Calculus, Vol. 2: Multi-Variable Calculus and Linear Algebra, with Applications to Differential Equations and Probability*. Blaisdell, Waltham, MA, 2nd ed. edition, 1969. [105]

[10] D. Attali and J.-D. Boissonnat. A linear bound on the complexity of the Delaunay triangulation of points on polyhedral surfaces. In *Proceedings of the seventh ACM Symposium on Solid modeling and Applications*, pages 139–146. ACM Press, 2002. [91]

[11] D. Attali, J.-D. Boissonnat, and H. Edelsbrunner. Stability and computation of medial axes — a state-of-the-art report. *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, To appear. [23]

[12] F. Aurenhammer. Voronoi diagrams – a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23(3):345–405, 1991. Habilitationsschrift. [Report B 90-09, FU Berlin, Germany, 1990]. [18]

[13] F. Aurenhammer and R. Klein. Voronoi diagrams. In J. Sack and G. Urrutia, editors, *Handbook of Computational Geometry, Chapter V*, pages 201–290. Elsevier Science Publishing, 2000. [SFB Report F003-092, TU Graz, Austria, 1996]. [18]

[14] C. Bajaj, H.Y. Lee, R. Merkert, and V. Pascucci. Nurbs based b-rep models for macromolecules and their properties. In *Proceedings of the fourth ACM Symposium on Solid modeling and Applications*, pages 217–228. ACM Press, 1997. [9, 76]

[15] C.L. Bajaj, V. Pascucci, A. Shamir, R.J. Holt, and A.N. Netravali. Dynamic maintenance and visualization of molecular surfaces. *Discrete Applied Mathematics*, 127(1):23–51, 2003. [9, 76]

[16] D.P. Bertsekas. *Convex Analysis and Optimization*. Athena Scientific, Belmont, MA, 2003. [105]

[17] J. Bloomenthal. Polygonization of implicit surfaces. *Computer Aided Geometric Design*, 5(4):341–355, 1988. [75]

[18] J. Bloomenthal and B. Wyvill, editors. *Introduction to Implicit Surfaces*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1997. [1, 7]

[19] H. Blum. A transformation for extracting new descriptors of shape. In W. Wathen-Dunn, editor, *Models for the Perception of Speech and Visual Form*, pages 362–380. Cambridge, MA, 1967. M.I.T. Press, 1967. [23]

[20] J.-D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Computational Geometry: Theory and Applications*, 22(1-3):185–203, 2002. [11]

[21] J.-D. Boissonnat, D. Cohen-Steiner, and G. Vegter. Isotopic implicit surface meshing. In *ACM Symposium on Theory of Computing (STOC)*, pages 301–309, 2004. [76]

[22] J.-D. Boissonnat and M. Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998. Translated by Hervé Brönnimann. [91]

[23] D. Braess. *Finite elements: Theory, fast solvers, and applications in solid mechanics*. Cambridge University Press, 1997. [6]

[24] H. Brönnimann, C. Burnikel, and S. Pion. Interval arithmetic yields efficient dynamic filters for computational geometry. *Discrete Applied Mathematics*, 109:25–47, 2001. [95]

[25] J. Buchner and T. Kiefhaber, editors. *Protein Folding Handbook*. John Wiley & Sons, Inc., New York, NY, USA, 2005. [2]

[26] E. Catmull and J. Clark. Recursively generated b-spline surfaces on arbitrary topological meshes. *Computer-Aided Design*, 10(6):350–355, 1978. [11]

[27] H-L. Cheng, T. K. Dey, H. Edelsbrunner, and J. Sullivan. Dynamic skin triangulation. *Discrete and Computational Geometry*, 25:525–568, 2001. [76, 98]

[28] H.-L. Cheng, H. Edelsbrunner, and P. Fu. Shape space from deformation. *Computational Geometry: Theory and Applications*, 19:191–204, 2001. [1, 2, 26, 103]

[29] H.-L. Cheng and X. Shi. Guaranteed quality triangulation of molecular skin surfaces. In *VIS '04: Proceedings of the conference on Visualization '04*, pages 481–488, Washington, DC, USA, 2004. IEEE Computer Society. [26, 76, 98]

[30] H.-L. Cheng and X. Shi. Quality mesh generation for molecular skin surfaces using restricted union of balls. In *VIS '05: Proceedings of the conference on Visualization '05*, Washington, DC, USA, 2005. IEEE Computer Society. [76]

[31] H.-L. Cheng and T. Tan. Approximating polygonal objects by deformable smooth surfaces. In *Mathematical Foundations of Computer Science*, pages 248–259, 2005. [35]

[32] H.L. Cheng. *Algorithms for Smooth and Deformable Surfaces in 3D*. PhD thesis, University of Illinois at Urbana-Champaign, Urbana, Illinois, 2002. [26]

[33] S.-W. Cheng, H. Edelsbrunner, P. Fu, and K.-P. Lam. Design and analysis of planar shape deformation. *Computational Geometry: Theory and Applications*, 19:205–218, 2001. [1, 2, 26, 103]

[34] L. Chew. Guaranteed quality Delaunay meshing in 3d. In *Proceedings of the thirteenth Annual ACM Symposium on Computational Geometry*, pages 391–393, 1997. [6, 75, 91, 94]

[35] H.I. Choi, S.W. Choi, and H.P. Moon. Mathematical theory of medial axis transform. *Pacific Journal of Mathematics*, 181(1):57–88, November 1997. [23]

[36] S. Choi and N. Amenta. Delaunay triangulation programs on surface data. In *The 13th ACM-SIAM Symposium on Discrete Algorithms*, pages 135–136, 2002. [91]

[37] Computational Geometry Algorithms Library. `http://www.cgal.org`. [94]

[38] M.L. Connolly. Analytical molecular surface calculation. *Journal of Applied Crystallography*, 16(5):548–558, Oct 1983. [9, 76]

[39] M.L. Connolly. Molecular surface Triangulation. *Journal of Applied Crystallography*, 18(6):499–505, Dec 1985. [9]

[40] M.L. Connolly. Molecular surfaces: A review. Number 14 in Computational Chemistry. Network Science, April 1996. http://www.netsci.org/Science/Compchem/feature14.html. [8]

[41] O. Devillers, S. Meiser, and M. Teillaud. The space of spheres, a geometric tool to unify duality results on voronoi diagrams. Technical report, INRIA - Sophia Antipolis, 1992. [26, 77]

[42] Tamal K. Dey and Jian Sun. An adaptive mls surface for reconstruction with guarantees. In *Symposium on Geometry Processing*, pages 43–52, 2005. [11]

[43] T.K. Dey and S. Goswami. Tight cocone: a water-tight surface reconstructor. In *Proceedings of the eighth ACM Symposium on Solid modeling and Applications*, pages 127–134. ACM Press, 2003. [47]

[44] T.K. Dey, G. Li, and T. Ray. Polygonal surface remeshing with delaunay refinement. In *Proc.14th Internat. Meshing Roundtable*, pages 343–361, 2005. [75]

[45] T.K. Dey and W. Zhao. Approximate medial axis as a voronoi subcomplex. In *Proceedings of the seventh ACM Symposium on Solid modeling and Applications*, pages 356–366, New York, NY, USA, 2002. ACM Press. [2, 25]

[46] G. L. Dirichlet. Über die Reduktion der positiven quadratischen Formen mit drei unbestimmten ganzen Zahlen. *J. Reine Angew. Math.*, 40:209–227, 1850. [17]

[47] N. Dyn, D. Levin, and C. A. Micchelli. Using parameters to increase smoothness of curves and surfaces generated by subdivision. *Computer Aided Geometric Design*, 7(1-4):129–140, 1990. [11]

[48] H. Edelsbrunner. Deformable smooth surface design. *Discrete and Computational Geometry*, 21:87–115, 1999. [1, 2, 21, 25, 26, 31, 68, 69, 76, 77, 81, 103, 118, 122]

[49] H. Edelsbrunner. *Geometry and topology for mesh generation.* Cambridge University Press, New York, NY, USA, 2001. [1, 7, 20]

[50] H. Edelsbrunner and E.P. Mücke. Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms. *ACM Transactions on Graphics*, 9(1):66–104, 1990. [20, 78]

[51] H. Edelsbrunner and A. Üngör. Relaxed scheduling in dynamic skin triangulation. In *Proceedings of the Japan Conference on Discrete and Computational Geometry*, pages 135–151, 2002. [76]

[52] G. Farin. *Curves and surfaces for CAGD: a practical guide*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2002. [1, 7]

[53] J. Flötotto. *A coordinate system associated to a point cloud issued from a manifold: definition, properties and applications*. PhD thesis, University of Nice-Sophia Antipolis, 2003. [11]

[54] S. J. Fortune. Voronoi diagrams and delaunay triangulations. *Euclidean Geometry and Computers*, pages 193–233, 1992. World Scientific Publishing Co., D.A. Du, F.K. Hwang, eds. [18]

[55] S. Funke, C. Klein, K. Mehlhorn, and S. Schmitt. Controlled perturbation for delaunay triangulations. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1047–1056, 2005. [20]

[56] S.M. Gelston and D. Dutta. Boundary surface recovery from skeleton curves and surfaces. *Computer Aided Geometric Design*, 12(1):27–51, 1995. [7, 23]

[57] P. Giblin and S. Brassett. Local symmetry of plane curves. *American Mathematical Monthly*, 92(10):689–707, 1985. [23]

[58] P. Giblin and B.B. Kimia. A formal classification of 3d medial axis points and their local geometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):238–251, 2004. [23]

[59] P.J. Giblin and B. B. Kimia. On the local form of symmetry sets, and medial axes, and shocks in 3D. In *Proc. of CVPR*, pages 566–573. IEEE Computer Society, 2000. [23]

[60] J.E. Goodman and J. O'Rourke, editors. *Handbook of Discrete and Computational Geometry*. CRC Press LLC, Boca Raton, FL, 2nd edition, 2004. [1, 16, 18]

[61] D. Halperin and E. Leiserowitz. Controlled perturbation for arrangements of circles. *International Journal of Computational Geometry & Applications*, 14(4 & 5):277–310, 2004. Special issue, papers from SoCG 2003. [20]

[62] E. Hartmann. A marching method for the triangulation of surfaces. *The Visual Computer*, 14(3):95–108, 1998. ISSN 0178-2789. [75]

[63] D. Hearn and M.P. Baker. *Computer graphics (2nd ed.): C version*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1997. [1]

[64] M. Held. Vroni: an engineering approach to the reliable and efficient computation of voronoi diagrams of points and line segments. *Computational Geometry: Theory and Applications*, 18(2):95–123, 2001. [18]

[65] H. Hoppe. *Surface reconstruction from unorganized points.* PhD thesis, Department of Computer Science and Engineering, University of Washington, June 1994. [11]

[66] H. Hoppe, T. DeRose, T. Duchamp, M. Halstead, H. Jin, J. McDonald, J. Schweitzer, and W. Stuetzle. Piecewise smooth surface reconstruction. In *Proc. SIGGRAPH 94*, pages 295–302, 1994. [11]

[67] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Comput. Graphics*, 26(2):71–78, 1992. Proc. SIGGRAPH '92. [11]

[68] B. Joe. Construction of three-dimensional Delaunay triangulations using local transformations. *Computer Aided Geometric Design*, 8(2):123–142, May 1991. [41]

[69] L. Kobbelt. $\sqrt{3}$-subdivision. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 103–112. ACM Press/Addison-Wesley Publishing Co., 2000. [5, 11, 75, 93]

[70] R. Kolluri. Provably good moving least squares. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1008–1018, August 2005. [11]

[71] E. Kreyszig. *Advanced Engineering Mathematics.* John Wiley & Sons, New York, NY, USA, 8th edition, 1999. [61]

[72] N.G.H. Kruithof and G. Vegter. Approximation by skin surfaces. In *Proceedings of the eighth ACM Symposium on Solid modeling and Applications*, pages 86–95. ACM Press, 2003. [33]

[73] N.G.H. Kruithof and G. Vegter. Triangulating skin surfaces. Technical Report ECG-TR-244303-01, Rijks*universiteit* Groningen, 2003. [75]

[74] N.G.H. Kruithof and G. Vegter. Approximation by skin surfaces. *Computer-Aided Design*, 36:1075–1088, 2004. [33, 72]

[75] N.G.H. Kruithof and G. Vegter. Meshing skin surfaces with certified topology. In *Proceedings of the nineth International CAD/Graphics conference*, pages 287–292, 2005. [75]

[76] N.G.H. Kruithof and G. Vegter. Meshing skin surfaces with certified topology. *Computational Geometry: Theory and Applications*, 2006. To appear. [75]

[77] B. Lee and F.M. Richards. The interpretation of protein structures: estimation of static accessibility. *Journal of Molecular Modeling*, 55:379–400, 1971. [8]

[78] D. Levin. Mesh-independent surface interpolation. *Geometric Modeling for Scientific Visualization*, pages 37–49, 2003. [11]

[79] David Levin. The approximation power of moving least-squares. *Mathematics of Computation*, 67(224):1517–1531, 1998. [11]

[80] A. Lieutier. Any open bounded subset of $\mathbb{R}^n$ has the same homotopy type than its medial axis. In *Proceedings of the eighth ACM Symposium on Solid modeling and Applications*, pages 65–75. ACM Press, 2003. [23]

[81] C. Loop. Smooth subdivision surface based on triangles. Master's thesis, Department of Mathematics, University of Utah, USA, 1987. [11]

[82] W.E. Lorensen and H.E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 163–169, New York, NY, USA, 1987. ACM Press. [11, 75]

[83] N. Mayya and V.T. Rajan. Voronoi diagrams of polygons: a framework for shape representation. *J. Math. Imaging Vis.*, 6(4):355–378, 1996. [18]

[84] D. Pedoe. *Geometry, a comprehensive course.* Dover Publications, New York, 1970. [26, 77]

[85] S. Plantinga and G. Vegter. Isotopic approximation of implicit curves and surfaces. In *ACM SIGGRAPH Symposium on Geometry Processing*, pages 251–260, 2004. [76]

[86] M.J.D. Powell. The theory of radial basis function approximation in 1990. *Advances in Numerical Analysis*, Vol. II: Wavelets, Subdivision Algorithms and Radial Basis Functions:105–210, 1992. [12]

[87] F.M. Richards. Areas, volumes, packing, and protein structure. *Annual Review of Biophysics and Bioengineering*, 6:151–176, 1977. [9]

[88] A.W. Roberts and D.E. Varberg. *Convex Functions.* Academic Press, New York and London, 1973. [105]

[89] J.-R. Sack and J. Urrutia. *Handbook of computational geometry.* North-Holland Publishing Co., Amsterdam, The Netherlands, The Netherlands, 2000. [1]

[90] C. Shen, J.F. O'Brien, and J.R. Shewchuk. Interpolating and approximating implicit surfaces from polygon soup. In *Proceedings of ACM SIGGRAPH 2004*, pages 896–904. ACM Press, August 2004. [11]

[91] R. Sibson. A brief description of natural neighbour interpolation. In Vic Barnett, editor, *Interpreting multivariate data*, pages 21–36. John Wiley & Sons, Chichester, 1981. [11]

[92] G.M. Treece, R.W. Prager, and A.H. Gee. Regularised marching tetrahedra: improved iso-surface extraction. *Computers and Graphics*, 23(4):583–598, 1999. [86, 88, 97]

[93] G. Turk, H.Q. Dinh, J. O'Brien, and Gary Yngve. Implicit surfaces that interpolate. In *Shape Modelling International*, pages 62–71, 2001. [12]

[94] G. Turk and J.F. O'Brien. Modelling with implicit surfaces that interpolate. *ACM Transactions on Graphics*, 21(4):855–873, 2002. [12]

[95] P.J. Vermeer. *Medial axis transform to boundary representation conversion.* PhD thesis, Department of Computer Sciences, Purdue University, 1994. [7, 23]

[96] G. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques, deuxième memoire, recherche sur les parallelloèdres primitifs. *Journal für die Reine und Angewandte Mathematik*, 134:198–287, 1908. [17]

[97] D.F. Watson. *Contouring : a guide to the analysis and display of spatial data.* Computer Methods in the Geosciences. Elsevier Science Publishers B. V., Amsterdam, The Netherlands, The Netherlands, 1992. [11]

[98] F.-E. Wolter. Cut locus an medial axis in global shape interrogation and representation. Memorandum 92-2, MIT, Ocean Engineering Dept., Design Laboratory, Cambridge, MA, USA, December 1993. Revised version. [23]

[99] P. Yushkevich, P.T. Fletcher, S. Joshi, A. Thall, and S.M. Pizer. Continuous medial representations for geometric object modeling in 2d and 3d. *Image and Vision Computing*, 2003. [7]

[100] D. Zorin, P. Schröder, and W. Sweldens. Interpolating subdivision for meshes with arbitrary topology. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 189–192, New York, NY, USA, 1996. ACM Press. [11]

Samenvatting

Voor het ontwikkelen en ontwerpen van producten wordt tegenwoordig veelvuldig gebruik gemaakt van software-pakketten waarmee drie-dimensionale computermodellen ontworpen kunnen worden. Deze modellen worden gebruikt om het product te visualiseren voordat een fysiek prototype wordt gemaakt en voor het simuleren van testen. Bekende voorbeelden hiervan zijn te vinden in de auto- en luchtvaartindustrie, maar ook consumentenproducten worden met behulp van deze pakketten ontwikkeld en getest.

Het modelleren van de oppervlakken van deze producten kun je je voorstellen als het boetseren van een stuk klei. Hierbij is de klei een computermodel waaraan je met de muis kunt "trekken" en "duwen". De methodes die we beschrijven in dit proefschrift worden gebruikt voor het modeleren van moleculen en gaan uit van een andere beschrijving van de oppervlakken, namelijk door middel van een verzameling bollen die binnen het oppervlak liggen. Het Michelinmannetje Bibendum is hier een duidelijk voorbeeld van. Hij bestaat uit een aantal bollen die door middel van gladde stukken aan elkaar geplakt worden.

Een eerste benadering voor het oppervlak gedefinieerd door een verzameling bollen is de rand van de vereniging van de bollen. Deze rand vormt een goede benadering van het oppervlak als voor ieder punt op het oppervlak er een bol in de invoer verzameling is die voldoende dicht bij dit punt ligt, zie ook figuur 1.2. Een nadeel is dat het oppervlak een knik heeft op de overgang tussen twee bollen. Met de methodes die in dit proefschrift beschreven worden is het mogelijk oppervlakken te maken die dicht om de rand van de vereniging van de bollen heen liggen en gladde stukken over intersectiekrommen hebben, waardoor de oppervlakken glad worden. Een voorbeeld van deze stukken is te zien in figuur 1.1.

Voor het maken van animaties is het nodig om een oppervlak in en ander oppervlak om te vormen. Er bestaan verschillende methodes die een verzameling bollen in een andere verzameling bollen deformeert. Met behulp van deze methodes is het mogelijk animaties te maken met de oppervlakken die in dit proefschrift worden beschreven.

Het eerste hoofdstuk van dit proefschrift bevat de introductie waarin de globale indeling van de rest van het proefschrift wordt gegeven en gerelateerd onderzoek wordt beschreven. In hoofdstuk 2 worden algemene wiskundige begrippen

geïntroduceerd die worden gebruikt in de daaropvolgende hoofstukken.

Hoofdstukken 3 en 4 beschrijven twee methodes om de gladde overgangen tussen de bollen te maken. Hierbij vormen envelope surfaces (beschreven in hoofdstuk 4) een generalisatie van de benaderingsmethode met behulp van skin surfaces uit hoofdstuk 3.

Het model dat wordt uitgewerkt in hoofdstuk 3 is gebaseerd op skin surfaces die door H. Edelsbrunner in [48] zijn ontwikkeld voor het modelleren van moleculen. Een skin surface wordt gedefinieerd door een verzameling bollen en een shrink factor. Ieder molecule bestaat uit een aantal atomen waarbij iedere atoom gerepresenteerd wordt door een bol. Het centrum van de bol komt overeen met de positie van het atoom, en de radius (Van der Waals radius) hangt af van het soort atoom. De rand van de vereniging van de bollen vormt het Van der Waals-oppervlak en wordt gebruikt om moleculen te visualiseren. Het Van der Waals-oppervlak is de skin surface van de bollen waarbij de shrink factor gelijk is aan 1. Voor een lagere, positieve, shrink factor neemt de radius van de invoerbollen af (de bollen worden kleiner) en verschijnen er gladde overgangen tussen de bollen. De skin surface van een caffeine molecule is te zien in figuur A.1. De shrink factor is in dit geval gelijk aan 1/2.

Voor het modelleren met behulp van skin surfaces is het verkleinen van de bollen ongewensd. We willen namelijk een oppervlak maken dat dicht om de vereniging van de bollen ligt en niet om de vereniging van een verzameling kleinere bollen. Daarom vergroten we de bollen eerst zó dat het verkleinen van de bollen bij het construeren van de skin surface teniet wordt gedaan. Door deze combinatie van vergroten van de bollen en het construeren van de skin surface krijgen we een klasse van oppervlakken die dicht om de vereniging van de bollen heen ligt en gladde stukken over de intersectie krommen heeft. Door middel van de shrink factor is het mogelijk om de grootte van de interpolerende stukken tussen de bollen te beïnvloeden. Een voorbeeld van benaderingen met verschillende shrink-factor-waarden is te zien in figuur 1.3. Voor een speciaal gekozen interval van shrink factor waarden tonen we aan dat de skin surface de volgende eigenschappen in relatie tot de rand van de vereniging van de bollen heeft. Ten eerste zijn de oppervlakken topologisch equivalent, wat onder meer inhoudt dat de twee oppervlakken hetzelfde aantal componenten, tunnels en gaten hebben. Verder liggen de twee oppervlakken niet ver uit elkaar (in termen van de Hausdorff-afstand) en raakt iedere invoerbol, die de rand van de vereniging raakt, de skin surface. Deze eigenschappen geven aan dat de benadering met behulp van een skin surface niet veel verschilt van de rand van de vereniging van de bollen.

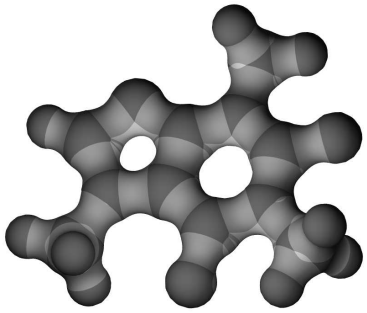

Figuur A.1: De skin surface van een caffeine molecule.

In hoofdstuk 4 generaliseren we de klasse van skin surfaces tot envelope surfaces. Met envelope surfaces is het mogelijk om de grootte van de interpolerende stukken lokaal aan te passen. Hierdoor kunnen platte delen van het oppervlak benaderd

worden met grote interpolerende stukken en kunnen kleinere stukken gebruikt worden op plaatsen waar het oppervlak meer detail heeft. Het nadeel van deze grotere flexibiliteit is dat een aantal eigenschappen die wel gelden voor skin surfaces niet per definitie gegarandeerd zijn voor envelope surfaces. Zo kan het oppervlak zelfdoorsnijdingen hebben of van topologie veranderen. We analyseren, met behulp van ideeën uit de convexiteitstheorie, onder welke condities de oppervlakken geen zelfdoorsnijdingen hebben en glad zijn.

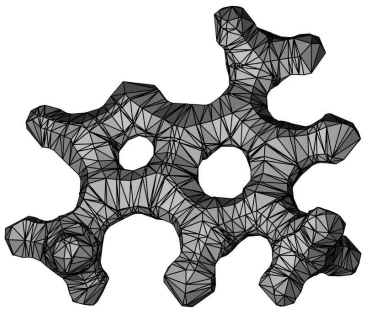

Figuur A.2: Een triangulatie van de skin surface van een caffeine molecule.

Tenslotte beschrijven we in hoofdstuk 5 een methode om een skin surface te trianguleren (meshen). Een triangulatie is een verzameling van driehoeken waarbij iedere driehoek een klein stukje van het oppervlak bedekt en de driehoeken op de randen aan elkaar passen. Een triangulatie van de skin surface van het caffeine molecule op de vorige pagina is te zien in de figuur aan de linker kant. Triangulaties zijn vaak nodig voor het visualiseren van de oppervlakken en voor verdere bewerkingen, zoals simulaties. Het triangulatie-algoritme maakt eerst een grove triangulatie die isotoop is met de skin surface. Dit houdt in dat de triangulatie en het skin surface topologisch equivalent zijn. De grove triangulatie kan grote driehoeken of langgerekte driehoeken bevatten. Voor visualisatie of voor verdere bewerkingen op de triangulatie moeten deze driehoeken worden vervangen. Dit kan worden gedaan door de triangulatie lokaal te verfijnen met behulp van post-processing algoritmes. Momenteel wordt dit algoritme geïmplementeerd en openbaar gemaakt als een uitbreiding van de Computational Geometry Algorithms Library (CGAL). We beschrijven deze implementatie kort en geven aan hoe dit algoritme gegeneraliseerd kan worden zodat het ook geschikt is om envelope surfaces te trianguleren.

Abstract

Nowadays, computer programs are often used for the construction of three dimensional models in the development and design of products. These models are used for visualization of the product even before the first prototype is constructed and for simulating tests. Well known examples can be found in the automotive- and aviation-industry. Also many consumer products are developed and tested with the use of these computer programs.

Generally, the surfaces of these models are developed by virtually modeling a piece of clay until it has the proper shape. In this analogue, the clay is a computer model that is deformable with the use of a mouse by "pushing" and "pulling" parts of the model. The methods described in this thesis use a different representation of the surface and were initially developed for modeling molecules. These surfaces are defined by a set of balls contained inside the surface. An example is the Michelin man "Bibendum", which is formed by a set of balls glued together by smooth patches.

A first approximation of the surface defined by a set of balls is the boundary of the union of these balls. This boundary forms a good approximation of the surface if there is a ball near every point on the surface, see also Figure 1.2. One of the disadvantages of this surface is that has cusps on the intersection curves where more than one ball touches the boundary. The methods described in this thesis are able to construct surfaces that wrap tightly around the union of the balls and have smooth patches over the intersection curves making the surface tangent continuous. See Figure 1.1 for an example of these patches.

For animations it is necessary to deform one surface into another. Several methods exist to deform one set of balls into another. These methods can also be used to deform the surfaces described in this thesis, hence giving an automatic way to deform between the surfaces.

The outline of this thesis is as follows. Chapter 1 contains the introduction in which related work is described and a more detailed outline of the thesis is given. In Chapter 2 common mathematical notions are introduced which are used in the subsequent chapters.

Chapters 3 and 4 describe two methods to construct tangent continuous surfaces that wrap tightly around the union of a set of balls and have smooth patches over the intersection curves. The envelope surfaces (described in Chapter 4) form a

generalization of the approximation method that uses skin surfaces from Chapter 3.

The model presented in Chapter 3 is based on skin surfaces, which are introduced by Edelsbrunner in [48] for the visualization of molecules. A skin surface is defined by a set of balls and a shrink factor. Each molecule consists of a set of atoms and each atom is represented as a ball centered on the location of the atom and its radius (Van der Waals-radius) depends on the type of atom. The boundary of the union of these balls forms the Van der Waals surface and is used for visualization of the molecule. The Van der Waals surface is equal to the skin surface with a shrink factor equal to one. For a lower, positive, shrink factor the radius of the input balls decreases and smooth patches appear between the balls. The skin surface of a caffeine molecule with a shrink factor equal to $1/2$ is depicted in Figure A.1

In molecular biology, the decrease in the radius of the input balls is not important. However for modeling, the skin surface should wrap tightly around the union of the input balls and not around shrunk input balls. Therefore we first increase the radius (grow the input balls) in such a way that due to the shrinking of the balls by the construction of the skin surface, the balls obtain their original size again. By combining these two steps (growing the balls and constructing the skin surface) we obtain a class of tangent continuous surfaces that wrap tightly around the union of the balls. With the shrink factor we are able to control the size of the interpolating patches, as is seen in Figure 1.3. We are able to derive properties between the boundary of the union of the balls and the skin surface for a specially chosen interval of shrink factors. First, the two surfaces are homeomorphic, which, among other things, means that they have the same number of connected components, tunnels and holes. Secondly, the two surfaces lie close to each other (in terms of the Hausdorff distance) and every input ball touching the boundary of the union of the input balls also touches the skin surface. From these properties, we derive that the two surfaces do not differ too much.

In Chapter 4 we generalize the class of skin surfaces to envelope surfaces. Using envelope surfaces we are able to control the size of the interpolating patches locally. This allows us to use larger interpolating patches on parts where the surface is flat and smaller patches on parts of the surface with more detail. The disadvantage of the increased flexibility is that properties which are valid for skin surfaces are not necessarily valid for envelope surfaces. For example, an envelope surface can have self intersections or can change topology. With results from convexity theory we analyze under which conditions envelope surfaces are tangent continuous and do not have self intersections.

Finally, in Chapter 5 we describe a method to triangulate (mesh) a skin surface. A triangulation of a smooth surface is a set of triangles such that each triangle covers a small part of the surface and the triangles fit nicely together on their boundaries. A triangulation of the caffeine molecule can be seen in Figure A.2. It is often necessary to construct a triangulation of a surface for visualization and further geometric processing. The triangulation algorithm first constructs a coarse isotopic triangulation. This implies that the skin surface and the triangulation are homeomorphic. The coarse triangulation can still contain large and skinny triangles which are not suitable for further geometric processing. These triangles can be removed

by a post-processing step which refines these triangles. This algorithm is currently being implemented as an extension package to the Computational Geometry Algorithms Library (CGAL). We briefly describe the implementation and the adaptations needed to make it suitable for meshing envelope surfaces.

Acknowledgements

Aan het einde van dit proefschrift wil ik een aantal mensen bedanken, zonder wie dit proefschrift niet tot stand had kunnen komen.

Ten eerste wil ik mijn promotor Gert Vegter bedanken. Jij bood mij een AIO-plaats aan halverwege mijn afstudeerproject. Zonder dat aanbod zou ik een carrière in het onderzoek nooit hebben overwogen en een heleboel zijn misgelopen. Verder wil ik je bedanken voor het actief meedenken met mijn onderzoek en de vele inspirerende discussies. Je hebt me niet alleen geleerd om onderzoek te doen, maar ook om de resultaten helder op te schrijven.

I would like to thank the reading committee, Prof.dr. Tamal Dey, Prof.dr. Jean-Daniel Boissonnat and Prof.dr. Mark Overmars for reading my thesis carefully and giving useful comments.

I would like to thank INRIA Sophia Antipolis, ETH Zürich and the Max Planck Institute in Saarbrücken for their hospitality during my visits in 2002, 2003 and 2005 and especially Joachim Giessen and Lutz Kettner for working with me.

De lunch was altijd een heerlijk moment om te ontspannen. Ik bedank Amit Chattopadhyay, Geert Fekken, Sunayana Ghosh, Joost le Feber, Gert-Jan van der Heiden, Joop Helder, Theresa Helmholt-Kleefsman, Barteld Kooi, Dirk-Jan Kort, Erwin Loots, Arie de Niet, Simon Plantinga, Gerk Rozema, Ena Tiesinga en Rik Wemmenhove voor de vele gesprekken over van alles behalve onderzoek.

Mijn kamergenoten Hendrik Wietze de Haan, Gao Hui, Simon Plantinga en Adrian Sandovici bedank ik voor de gezellige pauzes en de (over het algemeen sterke) koffie. Met name Hendrik Wietze en Simon wil ik bedanken voor de opvoeding in klassieke muziek. Ik ben geen natuurtalent, maar heb er wel erg van genoten.

Verder wil ik alle vrienden en kennissen bedanken voor de afleiding van en de interesse in de voortgang van mijn onderzoek. Met name: Corné & Annette, Marlies, Jacob & Baukje, Geke & Ewout, Ellen & Matthijs en Alexander & Gezien.

Tenslotte had ik dit proefschrift nooit kunnen schrijven zonder de continue steun en betrokkenheid van Irene en mijn (schoon-) ouders, André, Karin & Patrick, Mariëlle en René, Robert en Rieneke en Marten. Mariëlle, ik ben benieuwd wie het eerst "bevalt". Irene, bedankt voor je luisterende oor en je hulp waar dat ook maar mogelijk was. Ik zal nooit meer vergeten om bij een resultaat te vertellen hoeveel maanden ik er mee bezig ben geweest.