

Minkowski sum boundary surfaces of 3D-objects

Martin Peternell Helmut Pottmann Tibor Steiner

August 2, 2005

Institute of Discrete Mathematics and Geometry,
University of Technology Vienna, Austria

Abstract

Given two solids A and B with piecewise smooth boundary we discuss the computation of the boundary Γ of the *Minkowski sum* $A + B$. This boundary surface Γ is part of the envelope when B is moved by translations defined by vectors $\mathbf{a} \in A$, or vice versa. We present an efficient algorithm working for dense point clouds or for triangular meshes. Besides this, the global self intersections of the boundary Γ are detected and resolved. Additionally we point to some relations between Minkowski sums and kinematics, and compute local quadratic approximations of the envelope.

keywords: Minkowski sum, convolution surface, translation, motion, envelope, marching algorithm, point-set surface, signed distance function.

1 Introduction

Minkowski sums in two and three dimensions are used in various fields of mathematics, e.g., mathematical morphology, computer graphics, convex geometry and computational geometry. Given two sets A and B , the set

$$\{\mathbf{a} + \mathbf{B} \mid \mathbf{a} \in A\} = \{\mathbf{a} + \mathbf{b} \mid \mathbf{a} \in A, \mathbf{b} \in B\} = A + B \quad (1)$$

is called the *Minkowski sum* $A + B$ of the sets A, B . It is formed by translating the set B by all vectors $\mathbf{a} \in A$, and it is obvious that the sets A and B can be exchanged. The definition does not require any smoothness, but later on we will concentrate on sets A, B with piecewise smooth boundaries ∂A and ∂B .

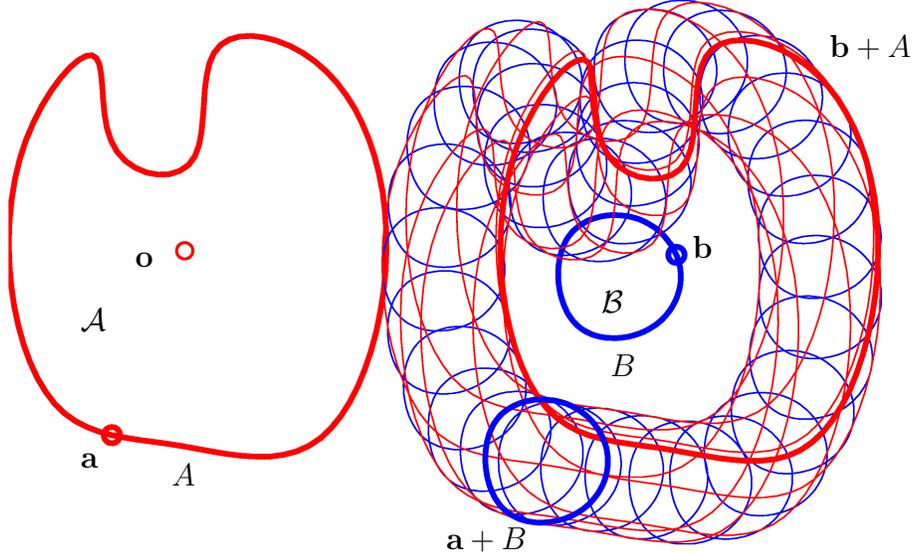


Figure 1: Construction of two families of curves $\mathbf{a} + B$ and $\mathbf{b} + A$

Since $A + B$ is generated by applying translations

$$\mathbf{x}' = \mathbf{a} + \mathbf{x}, \quad \mathbf{a} \in A, \mathbf{x}, \mathbf{x}' \in \mathbb{R}^3.$$

to the set B , the boundary $\partial(A + B)$ is part of the *envelope* of B with respect to these translations. In general we prefer to use boundary representations for geometric objects in the plane as well as in space. The Minkowski sum shall also be represented by its boundary surface, it is sufficient to apply translations defined by vectors $\mathbf{a} \in \partial A$ to the boundary points of B . When studying envelopes it is not necessary that the input objects are solids but can be surfaces as well.

Applying the concept of Minkowski sum to surfaces A and B may cause some confusion, since the boundary $\partial(A + B)$ of the sum of two surfaces differs from the sum $\partial A + \partial B$.

Figure 1 shows the situation for two planar curves A, B bounding planar domains \mathcal{A}, \mathcal{B} . The set $A + B$ can be covered by two families of curves. We choose an $\mathbf{a} \in A$ and form $\mathbf{a} + B$. If this is performed for all points \mathbf{a} , we obtain a family of congruent curves which are all translates of each other. In the same way the area is covered by the curves $\mathbf{b} + A$, for a moving point $\mathbf{b} \in B$. These two families of curves $\mathbf{a} + B$ and $\mathbf{b} + A$ possess an envelope E which might consist of several real components. Here we are mainly interested to compute the outer boundary Γ , the trimmed component of the envelope. This outer boundary Γ coincides with the boundary $\partial(\mathcal{A} + \mathcal{B})$ of the Minkowski sum of the sets \mathcal{A}, \mathcal{B} .

1.1 Previous work

There exist several contributions concerning the analysis of algorithms and the computation of Minkowski sums. In particular the 2D case is investigated thoroughly, for polygons as well as for smooth curves, see for instance [3, 5, 6, 7, 8].

The computation of the Minkowski sum of two convex objects, like planar polygons or polytopes, is known, see [10]. The handling of the non convex case is thus often solved by decomposing the input polytopes into their convex components or parts, see e.g. [4]. Then for all pairs of convex parts the Minkowski sum is computed. Finally one has to build the union of all these partial Minkowski sums. A convincing contribution is given by [16] and they can handle even complex objects. Additionally, the convex decomposition serves to estimate the combinatorial complexity. This technique is motivated by the computational or *constructive solid geometry* (CSG) representation of geometric objects, but when dealing with *boundary representations* it is not necessarily the appropriate tool.

If the input objects are represented by their boundaries, there are mainly two cases: The objects are given explicitly by parametric representations or the objects can be given by a piecewise planar approximation, consisting of triangular or polygonal faces. In general, the Minkowski sum computation with respect to parametric representations of the input objects results in a difficult reparametrization problem and is in close relation to the convolution of two surfaces, which is a superset of the final Minkowski sum, see [7, 9, 11, 14, 15].

We assume that the boundary of the input objects is given by a dense cloud of data points. This representation is on the one hand related to the parametric case when concerning local properties as normal vector estimation and local quadratic approximations. On the other hand, the Minkowski sum computation makes it necessary to take care of global phenomena. These can be handled with solid geometry representations or implicit representations like signed distance functions.

1.2 Contribution of the article

This article presents a technique to compute the Minkowski sum of two solids \mathcal{A}, \mathcal{B} with smooth or piecewise smooth boundaries. The boundary surfaces $A = \partial\mathcal{A}$, $B = \partial\mathcal{B}$ are considered to be represented by dense point sets or triangulated data, in order to estimate normal vectors for smooth parts and to detect edges and vertices in a stable way.

Under these assumptions, the outer boundary Γ of the Minkowski sum $A + B$ can be computed as follows: We consider A as set of translations $\mathbf{x}' = \mathbf{a} + \mathbf{x}$, $\mathbf{a} \in A$. This two-parameter family of motions generates a two-parameter family of surfaces $\mathbf{a} + B$ whose envelope is called the *convolution* $A \star B$ of A and B . The outer boundary of the envelope $A \star B$ is the boundary of the Minkowski sum $\mathcal{A} + \mathcal{B}$ of the sets \mathcal{A}, \mathcal{B} . The role of surface

A and B can be exchanged, and we can also compute the envelope of the two-parameter family $\mathbf{b} + A$ for $\mathbf{b} \in B$.

In most applications A and B are oriented surfaces and thus we can speak of an oriented envelope. Since A, B are defined by points \mathbf{a}, \mathbf{b} plus oriented unit normal vectors $\mathbf{n}_a, \mathbf{n}_b$, the same description will be available for the envelope $A \star B$ and the outer boundary Γ of the envelope. To extract the outer boundary Γ from the envelope $A \star B$, all parts lying inside have to be trimmed. This trimming is achieved by representing Γ by a *signed distance function*. The computation of this implicit representation is actually performed by evaluating the unsigned distance to $A \star B$ on points of a spatial grid and performing a marching algorithm on this grid to extract outside from inside points. Finally, these surface normal vectors can be used to detect and reconstruct global self intersections of Γ with relatively high reliability.

Further we discuss local quadratic approximations of the convolution surface $A \star B$. This serves to estimate the point density on the envelope depending on the point densities of the input surfaces A and B . Since the final outer boundary Γ is obtained by a trimming procedure of $A \star B$, where in fact the size of the chosen grid plays some role, it is important to guarantee the point density on Γ .

This paper is organized as follows. Section 2 is an introduction to convolution surfaces and envelope computation. The case of parametric input surfaces as well as the local quadratic approximations are briefly discussed in Section 3. Section 4 tells about convolution of point clouds, extraction of the outer boundary and the detection of self intersections. Additionally we give some remarks on the implementation.

2 Geometric background

Let \mathcal{A}, \mathcal{B} be two solids in space with smooth or piecewise smooth boundary surfaces $A = \partial\mathcal{A}, B = \partial\mathcal{B}$. We discuss the computation and some geometric properties of the Minkowski sum $\mathcal{A} + \mathcal{B}$. Since we focus on boundary representations of objects, we are interested in a mathematical description of $\partial(\mathcal{A} + \mathcal{B})$.

2.1 Convolution of two surfaces

Let us consider two piecewise smooth surfaces A, B . We identify points $\mathbf{a} \in A$ with translations $\mathbf{x}' = \mathbf{a} + \mathbf{x}$. If \mathbf{a} varies at the surface A we obtain a two-parametric family of translations. We can assume that there exists a (local) parametrization $\mathbf{a}(u, v)$ of A . At any smooth position \mathbf{a} there exist linearly independent partial derivative vectors

$$\mathbf{a}_u = \frac{\partial \mathbf{a}}{\partial u}, \mathbf{a}_v = \frac{\partial \mathbf{a}}{\partial v}.$$

Let $\mathbf{a}(t) = \mathbf{a}(u(t), v(t))$ be a curve in A and $\mathbf{x}' = \mathbf{a}(t) + \mathbf{x}$ be the corresponding one-parameter family of translations. The velocity vector field $\mathbf{v}(t) = \dot{\mathbf{a}}(t) = \mathbf{a}_u u_t + \mathbf{a}_v v_t$ is independent of the point \mathbf{x} and $\mathbf{v}(t_0)$ is a constant vector at any time instant t_0 .

For computational reasons and since we are interested in boundary representations we do not apply all translations $\mathbf{a} \in A$ to all points $\mathbf{b} \in B$ but we are looking for those pairs of points \mathbf{a}, \mathbf{b} such that $\mathbf{a} + \mathbf{b}$ is a boundary point of the envelope of the two-parameter family of surfaces $\mathbf{a} + B$.

For the given family of motions $\mathbf{x}' = \mathbf{a}(u, v) + \mathbf{x}$ only those points $\mathbf{b} \in B$ will contribute to the envelope, if both \mathbf{a}_u and \mathbf{a}_v are tangent to the surface B at the point \mathbf{b} . This in fact requires smoothness of the input surfaces A and B . In other words, $\mathbf{a} + \mathbf{b}$ is a point of the envelope if the tangent planes T_a and T_b at points \mathbf{a}, \mathbf{b} are parallel. Since we assume that A, B are oriented surfaces it is more precise to formulate a condition involving oriented unit normal vectors $\mathbf{n}_a, \mathbf{n}_b$ of A, B . This leads to the notion of *convolution surfaces* which are defined in the following way:

Given two oriented smooth surfaces A and B with unit normal vectors $\mathbf{n}_a, \mathbf{n}_b$, the convolution surface $A \star B$ is defined by

$$A \star B = \{\mathbf{a} + \mathbf{b} | \mathbf{n}_a = \mathbf{n}_b, \mathbf{a} \in A, \mathbf{b} \in B\}. \quad (2)$$

Since $A \star B$ is defined with respect to orientations of the normal vectors, it is an *oriented envelope* of B with respect to the family of motions. If A, B are smooth boundaries of sets \mathcal{A}, \mathcal{B} , the boundary of the Minkowski sum $\partial(\mathcal{A} + \mathcal{B})$ is usually a subset of the convolution $A \star B$,

$$\partial(\mathcal{A} + \mathcal{B}) \subset A \star B.$$

To figure out which points of the convolution $A \star B$ are actually boundary points of the Minkowski sum $\mathcal{A} + \mathcal{B}$, we perform a trimming of the convolution $A \star B$, see Figures 4 and 5.

2.2 Convolution of a surface and a curve

One can generalize the concept of Minkowski sum and convolution surfaces to the case where A is a smooth curve and B is a piecewise smooth surface in a straightforward manner. We consider B as boundary of a solid \mathcal{B} and let $\mathbf{a}(t)$ be a parametrization of A .

The convolution is then defined by

$$C = A \star B = \{\mathbf{a} + \mathbf{b} | \dot{\mathbf{a}} \cdot \mathbf{n}_b = 0, \mathbf{a} \in A, \mathbf{b} \in B\}. \quad (3)$$

The vector sum is formed for those pairs of points where the tangent vector $\dot{\mathbf{a}}$ is parallel to the tangent plane of \mathbf{b} . Since the normal vectors \mathbf{n}_a of the curve A are perpendicular to the tangent vectors $\dot{\mathbf{a}}$, this definition is equivalent to (2).

The convolution is the envelope of B with respect to a one-parameter family of motions $\mathbf{x}' = \mathbf{a}(t) + \mathbf{x}$ whose instantaneous velocity vector field equals $\dot{\mathbf{a}}(t)$. Thus, the computation of the convolution of a curve and a surface is simpler than for two surfaces. This concept applies to the convolution of a surface and a pipe surface which can be considered as offset of its center curve. Moreover it is useful when computing convolution surfaces of piecewise smooth objects which possess sharp edges, see Section 4.2.

3 Computing the envelope for parametrized surfaces

Given two smooth oriented surfaces $A, B \subset \mathbb{R}^3$ by parametric representations

$$\mathbf{a}(u, v) : G \rightarrow A, \quad \mathbf{b}(s, t) : H \rightarrow B,$$

where G and H are planar domains in \mathbb{R}^2 . The unit normal vectors of A, B are denoted by $\mathbf{n}_a(u, v)$ and $\mathbf{n}_b(s, t)$. The normalization of the normal vectors is not essential.

In order to compute the envelope of B with respect to translations defined by \mathbf{a} , or equivalently, the convolution surface $A \star B$, we have to construct a reparametrization

$$\begin{aligned} \phi : H &\rightarrow G, \\ (s, t) &\mapsto (u, v) \end{aligned}$$

such that $\mathbf{n}_a(\phi^{-1}(u, v)) = \mathbf{n}_b(s, t)$ holds. The mapping ϕ is assumed to be a *local diffeomorphism*, but globally it needs not to be one-to-one. Only in special cases there exist explicit analytic representations of ϕ .

Under the assumption that we have found $\mathbf{a}(\phi^{-1}(u, v))$, $\mathbf{b}(s, t)$, the envelope $A \star B$ is parametrized by

$$\mathbf{c}(s, t) = \mathbf{a}(\phi^{-1}(u, v)) + \mathbf{b}(s, t) = \mathbf{a}(s, t) + \mathbf{b}(s, t).$$

In some cases the convolution of two surfaces can explicitly be computed. We like to note that the *convolution of two spheres* is a sphere, which is geometrically evident, since if the center of a sphere moves along another sphere, the envelope is again a sphere. An analogous result holds for two paraboloids, if their axes are parallel, see section 3.1.

Further, it is shown in [11] that the convolution of any rational surface and a paraboloid is again a rational surface, and the reparametrization can be given explicitly. The same property holds for a larger class of surfaces, whose normal vectors admit a linear parametrization in the surface parameters, see [14]. In this sense, these surfaces are generalizations of paraboloids. Convolution surfaces of two ruled surfaces can also be parametrized explicitly, see [9]. Similar properties hold for two canal surfaces. The existence of explicit parametrizations is in close relation with the structure of the unit normal vector fields of the input surfaces.

Moreover, we like to mention some *degenerate cases* concerning the unit normal vector field of the surfaces. Let A be planar, then \mathbf{n}_a is a constant vector. Assume that $\mathbf{n}_b(s, t)$ is an injective parametrization of part of the unit sphere S^2 , then there exists at most a single point $\mathbf{b}_0 = \mathbf{b}(s, t)$ whose normal vector satisfies $\mathbf{n}_b(s, t) = \mathbf{n}_a$. In other words, \mathbf{b}_0 corresponds to all points $\mathbf{a} \in A$. We conclude that $A \star B$ is planar with $\mathbf{a}(u, v) + \mathbf{b}_0$ as possible parametrization.

Let A be a developable surface, then $\mathbf{n}_a(u, v)$ is a curve in S^2 . Without loss of generality we assume this curve to be parametrized by $\mathbf{n}_a(u)$. This says in particular that $\mathbf{a}(u_0, v)$ parametrizes a generating line in A for fixed u_0 . If $\mathbf{n}_b(s, t)$ is injective, there exists in general a curve $\mathbf{b}(s(u), t(u))$ which contributes to the construction. Since all points $\mathbf{a}(u_0, v)$ correspond to a single point $\mathbf{b}(s(u_0), t(u_0))$, the sum $\mathbf{a}(u_0, v) + \mathbf{b}(s(u_0), t(u_0))$ is a generating line in $A \star B$. Points on this line possess unit normals $\mathbf{n}_a(u_0, v)$. Since this property holds for all $(u, v) \in G$, the envelope $A \star B$ is a developable surface, too.

Convolution surfaces of translated input objects

For practical considerations it is remarkable that *the convolution $A \star B$ is invariant with respect to translations* in the following sense: Let $\tau : \mathbf{x}' = \mathbf{v} + \mathbf{x}$ be a fixed translation, and let $B' = \tau(B)$. Thus, B' admits the parametrization $\mathbf{b}'(s, t) = \mathbf{b}(s, t) + \mathbf{v}$. We assume that $\mathbf{a}(s, t)$ is already a parametrization of A with respect to coinciding unit normals $\mathbf{n}_a(s, t) = \mathbf{n}_b(s, t)$. Because $\mathbf{b}(s, t)$ and $\mathbf{b}'(s, t)$ possess the same unit normal vectors \mathbf{n}_b , the convolution $A \star B'$ is parametrized by

$$\mathbf{a} + \mathbf{b}' = \mathbf{a} + \mathbf{b} + \mathbf{v} = \mathbf{c} + \mathbf{v}.$$

Thus, the convolution surface $A \star B'$ of A and B is obtained by applying the translation $\tau : \mathbf{x}' = \mathbf{v} + \mathbf{x}$ to the convolution surface $C = A \star B$ of A and B .

Convolution surfaces of offsets of the input objects

Convolution surfaces are *invariant with respect to the offset operation* in the following sense: Given two smooth and oriented surfaces A and B by parametrizations $\mathbf{a}(s, t)$ and $\mathbf{b}(s, t)$ where we assume that the reparametrization with respect to coinciding unit normal vectors $\mathbf{n}_a = \mathbf{n}_b$ has already performed.

How does the convolution surface look like, if we replace B by a one-sided offset B_d at oriented distance d . In fact, B and B_d have same unit normals \mathbf{n}_b , since $\mathbf{b}_d(s, t) = \mathbf{b}(s, t) + d\mathbf{n}_b(s, t)$. Since the unit normals of B, B_d and C coincide, we denote them by \mathbf{n} . As a direct consequence, the convolution $C' = A \star B_d$ is an offset of $C = A \star B$ at distance d , and a parametrization of C' is obtained by

$$\mathbf{c}' = \mathbf{a} + \mathbf{b}_d = \mathbf{a} + \mathbf{b} + d\mathbf{n} = \mathbf{c} + d\mathbf{n}.$$

Example: Let A be a piecewise smooth object and B be a *pipe surface* with *spine curve* \mathbf{m} and radius d . The pipe surface B itself is the envelope of a one-parameter family of spheres traveling along \mathbf{m} . We interpret $M = \mathbf{m}$ as offset of B at distance $-d$, and the unit normals \mathbf{n}_b of B are perpendicular to M (for corresponding parameter values). Thus, the convolution $C = A \star B$ can be computed as the offset of $C' = A \star M$ at oriented distance d . The surface C' is the convolution of a surface A and a curve M , see Section 2.2.

3.1 Convolution surfaces of spheres and paraboloids and the quadratic approximation of general convolution surfaces

To get some insight into the computation of convolution surfaces, we discuss the convolution of two spheres and two paraboloids with parallel axes. Later is a special case of the quadratic approximation of the convolution surface of two arbitrary analytic input surfaces. Figure 2 is an illustration of the planar case, showing the convolution of two circles and two parabolas with parallel axes.

Convolution of spheres

Let A and B be two oriented spheres with centers \mathbf{r} and \mathbf{s} . Their signed radii are denoted by r and s , and determine the orientation of A and B . The convolution $C = A \star B$ is again a sphere, with center $\mathbf{r} + \mathbf{s}$ and radius $r + s$.

This can be checked by using standard parametrizations

$$\mathbf{a} = \mathbf{r} + r\mathbf{n}(u, v), \mathbf{b} = \mathbf{s} + s\mathbf{n}(u, v)$$

for both A and B , where $\mathbf{n}(u, v)$ is a parametrization of the unit sphere. Since the correspondence between A and B is defined by equal surface parameters (u, v) , the convolution surface $A \star B$ is parametrized by

$$\mathbf{c} = \mathbf{a} + \mathbf{b} = \mathbf{r} + \mathbf{s} + (r + s)\mathbf{n}(u, v). \quad (4)$$

Quadratic approximation of the convolution surface of two smooth surfaces

The computation of convolution surfaces is based on a simple operation, the sum of vectors, and the question arises how local approximations of convolution surfaces look like. As a byproduct of this investigation we obtain the result that the convolution of two paraboloids with parallel axis is again a paraboloid.

Let A, B be two analytic surfaces. We assume that the regular points $\mathbf{a}_0, \mathbf{b}_0$ are a pair of corresponding points of A, B with respect to coinciding unit normals. Thus we are allowed

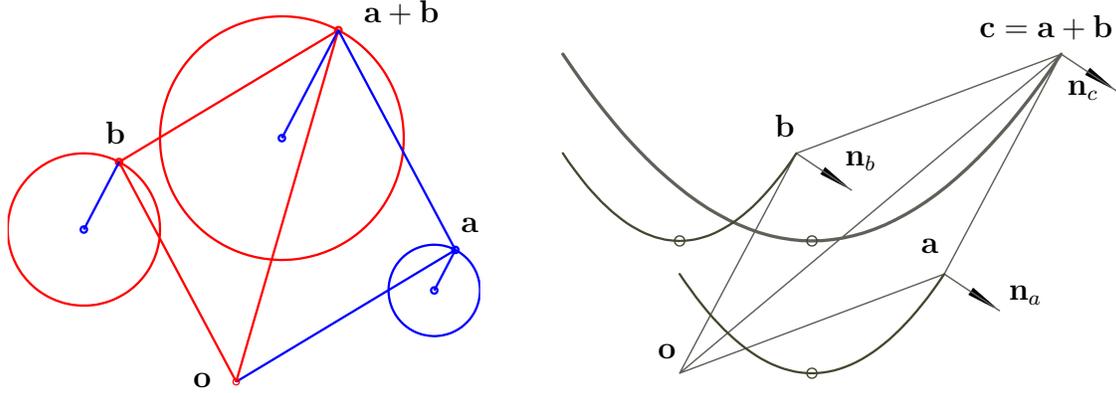


Figure 2: Convolution of two circles and two parabolas

to apply a translation to achieve $\mathbf{a}_0 = \mathbf{b}_0$. We choose a coordinate system with \mathbf{a}_0 as origin and in a way that A and B are locally parametrized by the following graph representations,

$$\mathbf{a}(u, v) = (u, v, au^2 + bv^2 + O(x^3)), \mathbf{b}(s, t) = (s, t, \alpha s^2 + \beta st + \gamma t^2 + O(x^3)), \quad (5)$$

where $O(x^n)$ denotes terms of order n . The third coordinate is expanded into a Taylor-series, and we do not take care of third or higher order terms. In order to obtain a compact and concise notation, we let $U = (u, v)^\top$, $S = (s, t)^\top$, and set

$$D = \begin{bmatrix} a & 0 \\ 0 & c \end{bmatrix}, \text{ and } E = \begin{bmatrix} \alpha & \beta \\ \beta & \gamma \end{bmatrix}.$$

By the use of these abbreviations, the local parametrizations of A, B are given by

$$\mathbf{a}(U) = \begin{bmatrix} U \\ U^\top \cdot D \cdot U + O(x^3) \end{bmatrix}, \text{ and } \mathbf{b}(S) = \begin{bmatrix} S \\ S^\top \cdot E \cdot S + O(x^3) \end{bmatrix}. \quad (6)$$

Since the gradients of the third coordinate a_3, b_3 of \mathbf{a}, \mathbf{b} equal

$$\nabla a_3 = 2D \cdot U + O(x^2), \quad \nabla b_3 = 2E \cdot S + O(x^2),$$

the reparametrization ϕ and its inverse according to coinciding unit normals are linearly approximated by the equations

$$\phi : U = D^{-1} \cdot E \cdot S + O(x^2), \quad \phi^{-1} : S = E^{-1} \cdot D \cdot U + O(x^2). \quad (7)$$

It is obvious that if \mathbf{a}_0 is a *parabolic point* of A , we have either $a = 0$ or $c = 0$, which implies that D is not invertible, and ϕ does not exist. If $\mathbf{b}_0 \in B$ is either elliptic or hyperbolic then a curve $c \in B$ with $\mathbf{b}_0 \in c$ contributes to this local construction. In the following we assume regularity of D and E .

According to the reparametrizations (7) we can either represent A with respect to S or B with respect to U ,

$$\mathbf{a}(S) = \begin{bmatrix} D^{-1} \cdot E \cdot S + O(x^2) \\ S^\top \cdot E \cdot D^{-1} \cdot E \cdot S + O(x^3) \end{bmatrix}, \text{ or } \mathbf{b}(U) = \begin{bmatrix} E^{-1} \cdot D \cdot U + O(x^2) \\ U^\top \cdot D \cdot E^{-1} \cdot D \cdot U + O(x^3) \end{bmatrix}, \quad (8)$$

and finally the convolution surface $A \star B$ admits the parametrization

$$(\mathbf{a} + \mathbf{b})(S) = \begin{bmatrix} (I + D^{-1} \cdot E) \cdot S + O(x^2) \\ S^\top \cdot (E + E \cdot D^{-1} \cdot E) \cdot S + O(x^3) \end{bmatrix}. \quad (9)$$

Analogously, $A + B$ can be represented with respect to U . In order to obtain a graph representation of the quadratic approximation of $A \star B$, we let $F = I + D^{-1} \cdot E$ and with the reparametrization

$$T = F \cdot S + O(x^2), \text{ or } S = F^{-1} \cdot T + O(x^2),$$

we finally get

$$(\mathbf{a} + \mathbf{b})(T) = \begin{bmatrix} T \\ T^\top \cdot F^{-1} \cdot E \cdot T + O(x^3) \end{bmatrix}. \quad (10)$$

As known from differential geometry, the principal curvatures of $A \star B$ are determined by the eigenvalues of $M = F^{-1} \cdot E$ and the associated directions are its eigenvectors. Moreover, the Gaussian and mean curvature of $A \star B$ equal $\det(M)$ and $\text{trace}(M)$.

Convolution of paraboloids with parallel axes

Equations (9) and (10) also prove that the convolution $A \star B$ of two paraboloids A, B with parallel axes is again a paraboloid. This is realized by consequently omitting all higher order terms which are given in brackets. More precisely, the parametrizations (6) of \mathbf{a} and \mathbf{b} represent paraboloids with parallel axes if we truncate the power series after second order terms. Further, since the reparametrization $\phi : U = D^{-1} \cdot E \cdot S$ is *linear*, $\mathbf{a}(S)$ (or $\mathbf{b}(U)$) are quadratic graph representations of A, B , too. Therefore, the parametrization of the convolution $A \star B$

$$(\mathbf{a} + \mathbf{b})(S) = \begin{bmatrix} (I + D^{-1} \cdot E) \cdot S \\ S^\top \cdot (E + E \cdot D^{-1} \cdot E) \cdot S \end{bmatrix}. \quad (11)$$

is a quadratic graph representation. It can be beautified by the final linear reparametrization $T = (I + D^{-1} \cdot E) \cdot S$ to obtain

$$(\mathbf{a} + \mathbf{b})(T) = \begin{bmatrix} T \\ T^\top \cdot F^{-1} \cdot E \cdot T \end{bmatrix}. \quad (12)$$

Point density on the convolution surface

Assuming that the input surfaces A, B are sampled with density δ , we compute the sampling density of the convolution $A \star B$ at regular surface points.

Consider a pair of corresponding regular surface points $\mathbf{a}_0, \mathbf{b}_0$ in A, B . By the choice of an appropriate coordinate system we achieve local parametrizations $\mathbf{a}(U), \mathbf{b}(S)$, as in (6).

Assume that $\|\mathbf{a}(S)\| < \delta$ and $\|\mathbf{b}(S)\| < \delta$ hold, then the triangle inequality leads to $\|(\mathbf{a} + \mathbf{b})(S)\| < 2\delta$. For the input surfaces we actually have

$$\begin{aligned}\|\mathbf{a}(S)\|^2 &= \|D^{-1} \cdot E \cdot S\|^2 + \|S^\top \cdot E \cdot D^{-1} \cdot E \cdot S\|^2 \leq \frac{1}{d^2} e^2 \sigma^2 (1 + \sigma^2 e^2), \\ \|\mathbf{b}(S)\|^2 &= \|S\|^2 + \|S^\top \cdot E \cdot S\|^2 \leq \sigma^2 (1 + \sigma^2 e^2)\end{aligned}$$

where $\sigma = \|S\|$, $e = \|E\|$ and $d = \|D\|$. In fact, $2d$ and $2e$ are the maximal principal curvatures of the input surfaces A and B at \mathbf{a} and \mathbf{b} . For simplicity we use the Euclidean norm for vectors and the largest eigenvalue (spectral norm) for matrices. Thus, the norm of $\mathbf{a} + \mathbf{b}$ is estimated by

$$\|(\mathbf{a} + \mathbf{b})(S)\| \leq \|\mathbf{a}(S)\| + \|\mathbf{b}(S)\| \leq \sigma \frac{e + d}{d} \sqrt{1 + \sigma^2 e^2}.$$

On the other hand we could estimate $\|\mathbf{a} + \mathbf{b}\|$ directly from (11), and obtain

$$\begin{aligned}\|(\mathbf{a} + \mathbf{b})(S)\|^2 &= \|(I + D^{-1} \cdot E) \cdot S\|^2 + \|S^\top \cdot (E + E \cdot D^{-1} \cdot E) \cdot S\|^2, \text{ thus} \\ \|(\mathbf{a} + \mathbf{b})(S)\| &\leq \sigma \frac{\sqrt{e^2 + d^2}}{d} \sqrt{1 + \sigma^2 e^2}.\end{aligned}$$

4 Envelopes of surfaces given by dense point clouds

Given two oriented surfaces A, B as sets of points $\mathbf{a}_i, i = 1, \dots, M$ and $\mathbf{b}_j, j = 1, \dots, N$ as well as their associated oriented unit normal vectors $\mathbf{n}_a, \mathbf{n}_b$. For a moment we assume that A, B are smooth surfaces and we will extend the construction to piecewise smooth surfaces later.

Two points \mathbf{a} and \mathbf{b} are said to be *corresponding*, if $\mathbf{n}_a = \mathbf{n}_b$ holds. For any \mathbf{a} there exists a set of corresponding points \mathbf{b}_j , which can also be empty. Consequently, the convolution $A \star B$ is formed by points $\mathbf{c}_j = \mathbf{a} + \mathbf{b}_j$. Usually, this correspondence through equal normal vectors is not one-to-one. Obviously, the surface normal vector \mathbf{n}_{c_j} of the convolution $A \star B$ at points \mathbf{c}_j equals \mathbf{n}_a .

Figure 3 shows on the left hand side the construction of corresponding points of the convolution $A \star B$ of two planar curves A, B , and the one-parameter family of translated curves $\mathbf{b} + A$ on the right hand side. Figure 4 shows the point set of the convolution $A \star B$ of two planar curves A, B on the left hand side and the extracted outer boundary Γ of the envelope $A \star B$ on the right hand side.

The practical computation is best performed not only on point clouds but on triangulated surfaces. Thus let A be given by data points \mathbf{a}_i and triangulation $T = \{t\}$. Analogously, let B be given by data points \mathbf{b}_j and triangulation $S = \{s\}$. To construct the correspondence

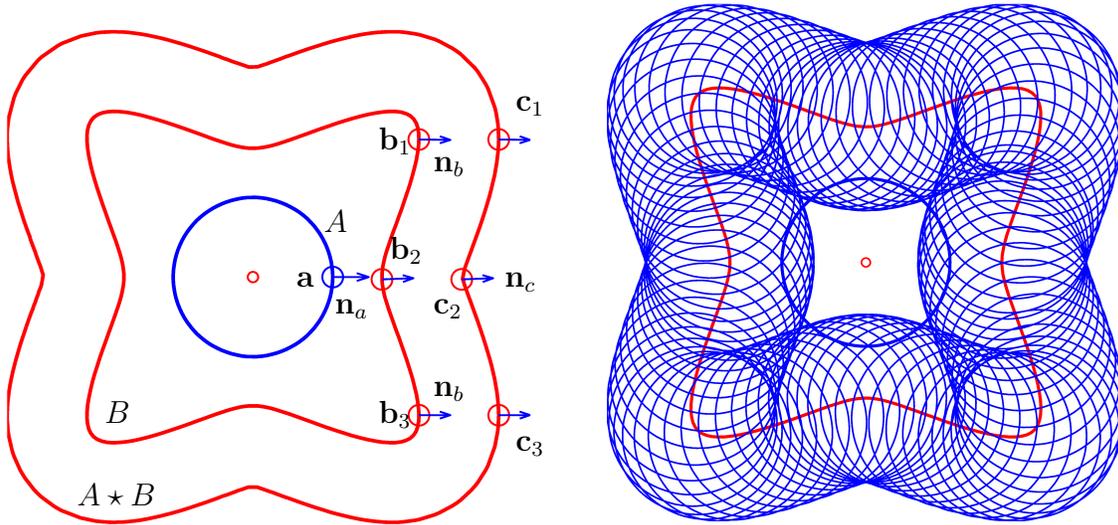


Figure 3: Corresponding points on A and B and family of translated curves $\mathbf{b} + A$

between A and B we proceed as follows: For each point \mathbf{a} with associated normal vector \mathbf{n}_a we search for those \mathbf{b}_j whose normal vectors \mathbf{n}_{b_j} satisfy

$$\|\mathbf{n}_a - \mathbf{n}_{b_j}\| < \tau, \text{ or } \angle(\mathbf{n}_a - \mathbf{n}_{b_j}) < \tau. \quad (13)$$

These two requirements are almost equivalent for small values τ . The set of triangles $S = \{s\}$ is mapped to a set of triangles $\{s'\}$ approximating the unit sphere. It is no longer a triangulation, but the set $\{s'\}$ can be used to find those triangles s' which contain the point \mathbf{n}_a . We mention that the notation 'contain' is not really correct since s' is not a spherical triangle. More precisely we would have to state that \mathbf{n}_a lies in the convex cone spanned by the vectors of the vertices of the triangle s' . But since s' is small and could be replaced by a spherical triangle with same vertices, we simply use the word 'contain'.

Let s' be a triangle in S^2 with vertices $\mathbf{n}_{b_k}, k = 1, 2, 3$ which contains \mathbf{n}_a . Then, \mathbf{n}_a can be represented by

$$\rho \mathbf{n}_a = \sum_{k=1}^3 \mathbf{n}_{b_k} \beta_k,$$

where ρ just denotes the norm of the right hand side vector. To compute the corresponding point to \mathbf{a} with respect to the chosen triangle, we use

$$\left(\frac{\beta_1}{\beta}, \frac{\beta_2}{\beta}, \frac{\beta_3}{\beta} \right), \text{ with } \beta = \sum_{k=1}^3 \beta_k$$

as barycentric coordinates within the 'original' triangle s with vertices \mathbf{b}_k on the input surface B . Consequently,

$$\mathbf{b} = \frac{1}{\beta} \sum_{k=1}^3 \mathbf{b}_k \beta_k$$

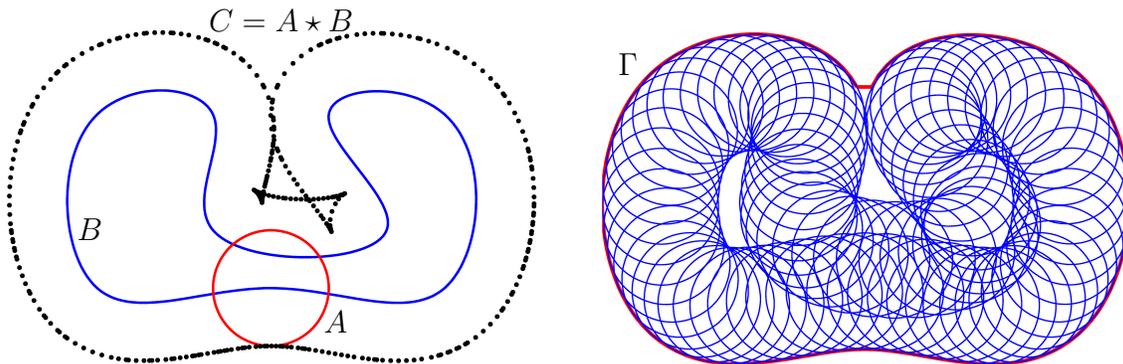


Figure 4: Envelope $C = A \star B$ and its outer boundary Γ .

corresponds to the given point \mathbf{a} . Performing this construction for each triangle s' containing \mathbf{n}_a , we are able to compute all points \mathbf{b}_j corresponding to the fixed point \mathbf{a} . Finally, $\mathbf{c}_j = \mathbf{a} + \mathbf{b}_j$ are points of the convolution surface $A \star B$, thus candidates for points of outer boundary of the Minkowski sum.

Remark: In order to define correspondences one can exchange \mathbf{a} by \mathbf{b} and perform the same construction. Additionally we note that if A and B are given by rather *dense point clouds*, the computationally costly linear interpolation within each triangle can be omitted. In this case it is sufficient that the correspondence is computed by searching for nearest neighbors of the sets of unit normal vectors $\mathbf{n}_a, \mathbf{n}_b$ according to (13).

So far we have computed an unorganized set of points \mathbf{c} with associated unit normals \mathbf{n}_c covering the convolution surface $C = A \star B$. We recall that C is the envelope of B with respect to translations determined by vectors $\mathbf{a} \in A$ or equivalently the envelope of A with respect to translations given by $\mathbf{b} \in B$.

4.1 Extracting the outer boundary from the convolution

Figure 5 shows that the envelope $C = A \star B$ might have global self intersections and possess parts in the interior. In practice these interior parts are of minor interest and one likes to obtain a representation of the outer boundary.

Local methods, often preferable, are unfortunately not an appropriate tool to detect global self intersections and to determine which parts of the envelope lie inside. Therefore we have to apply a global method to extract the outer boundary.

The result of the previous step is a point set \mathbf{c} with associated unit normal vectors \mathbf{n}_c of the envelope surface C . We do not care about holes inside the envelope and thus proceed as follows to compute the outer boundary surface.

The envelope C is placed in a rectangular grid whose cell size is larger than δ , the average

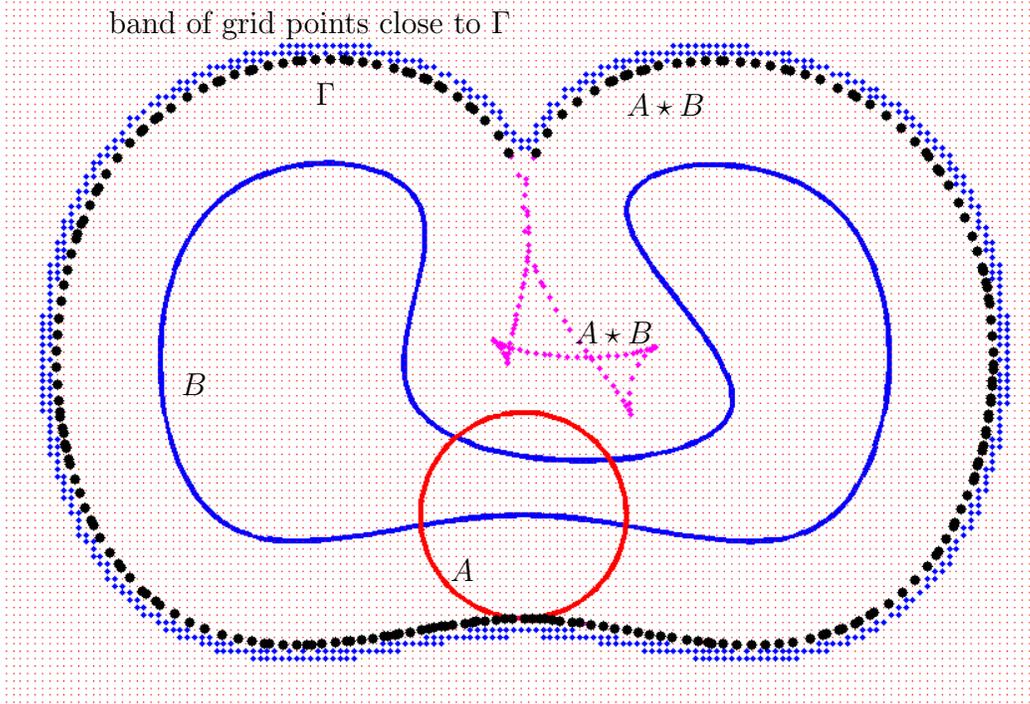


Figure 5: Extaction of the outer boundary Γ of the convolution $A \star B$.

point distance at the envelope. We compute the unsigned distance function of C , evaluated at the grid points. In order to trim those parts of C lying in the interior of the outer boundary Γ we perform a fast marching method similar to that described in [17, 12] on the grid points.

Roughly this works as follows: At the beginning all grid points are labeled as *unknown* and *inside*. We begin the loop by taking a point surely outside of Γ and initialize an *active front* by storing this point in it and label it as *trial*. We extract a grid point \mathbf{p} from the active front and this point \mathbf{p} is labeled as *visited* and *outside*. We collect all its *unknown* neighbors \mathbf{q} on the grid. If the distance value of \mathbf{q} is larger than some threshold $\tau(\sim 2.5\delta)$, \mathbf{q} is labeled as *outside* and *trial* and \mathbf{q} is added to the active front. In this manner the front propagates on the grid and this procedure is performed until the active front is empty.

The grid points are now correctly labeled as *inside* or *outside*. Since it is necessary to stop the front a little bit too early at $\sim 2.5\delta$, we let it move towards the envelope $A \star B$ by $\sim 2.0\delta$. This defines a *signed distance function* of the outer boundary Γ . By taking a small band of *outside* grid points, the outer boundary Γ is defined as set of nearest neighbors of this small band, see Figure 5. This results in a relatively uniform sampling of Γ . We note that also the normal vectors at points of Γ are available. The 2D case of the boundary extraction is illustrated in Figure 5. Some results in 3D are shown in Figure 6. Finally we mention that the extraction of the outer boundary can be significantly accelerated by considering only those grid points within a certain bandwidth ($\sim 5\delta$) of the convolution

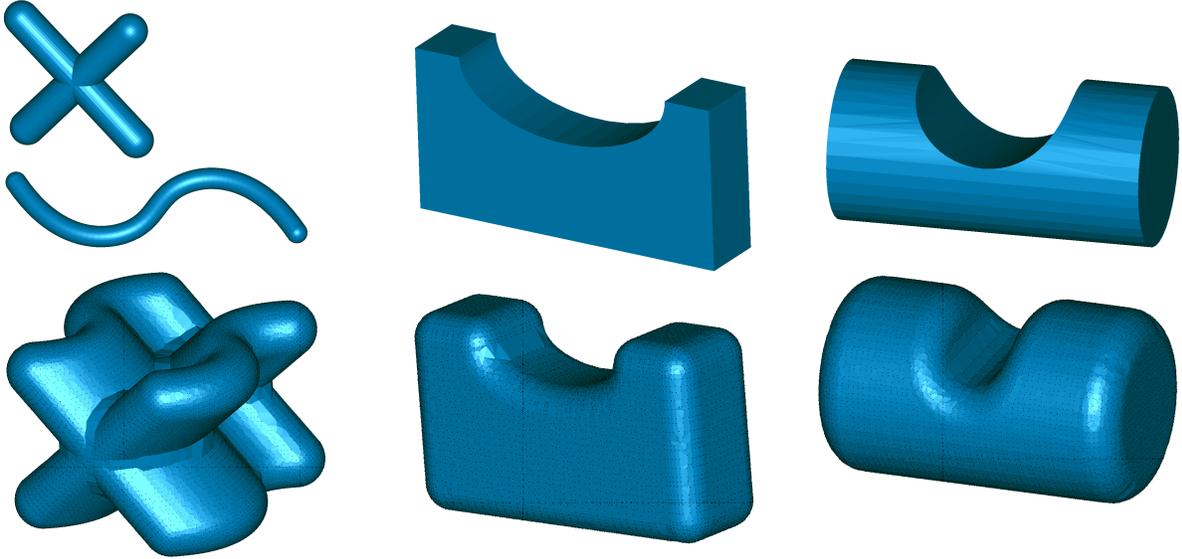


Figure 6: Left: Minkowski sum of two objects. Middle and right: Offset of cad-model computed via Minkowski sum.

C. Self intersection of the envelope Γ are smoothed out and are treated in a separate procedure.

4.2 Edges and corners

Real world objects might possess edges and corners and we have to extend the computation of convolution surfaces for this case. Singularities of surfaces from a differential or algebraic point of view might even be more complicated. For practical reasons we restrict ourselves here to *edges* and *vertices*.

In practice we assume that edges of the objects A, B can be reliably detected. If A, B are cad-models, the edges are computed by searching for adjacent faces f_1, f_2 (triangles) whose unit normals $\mathbf{n}_1, \mathbf{n}_2$ form an angle larger than some threshold τ . If A, B are represented by a dense point cloud from scan data, this criterion might be wrong because of measurement errors and noisy data. For those data, edges are more likely detected by searching for highly curved regions. But only those regions are possible edges which can be approximated sufficiently well by curves. Additionally it must be possible to define two different sides along an edge-like region.

At an edge point \mathbf{p} there exist two different unit normal vectors \mathbf{n}_l and \mathbf{n}_r at either side of the edge. We exclude curves of regression here. The set of normal vectors at \mathbf{p} is thus parametrized by

$$(1 - u)\mathbf{n}_l + u\mathbf{n}_r, \quad \text{with } u \in [0, 1].$$

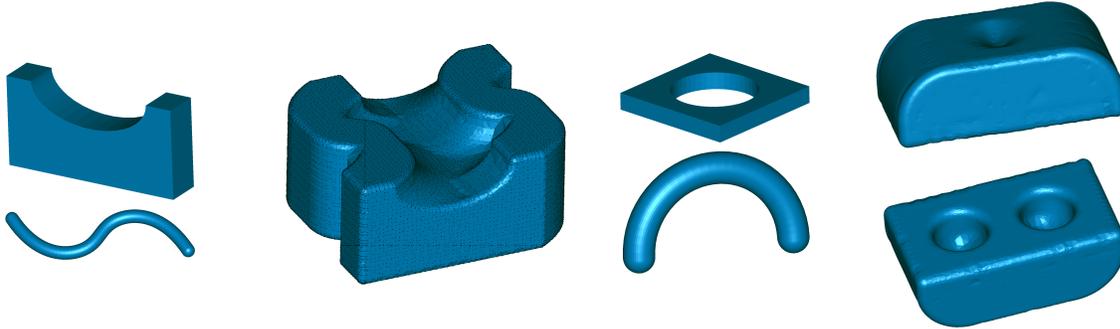


Figure 7: Left: Cad-model and pipe surface and their Minkowski sum. Right: Plate with hole and half torus and two views of their Minkowski sum.

The corresponding unit normals describe a circular arc.

At a corner \mathbf{p} the set of unit normal vectors is even two-dimensional and is defined by the convex cone of unit normals of the surface patches meeting at \mathbf{p} .

Since the set of relevant normals is larger at edges and corners, these points contribute usually more likely to the envelope. Thus it is important to detect sharp edges and corners at the objects to guarantee an accurate representation of the outer boundary. Figure 6 shows offsets of two cad models with precise edges and corners. These offsets are computed as Minkowski sums of the objects and a sphere. If one is only interested in offset computation, there are simpler algorithms since it is not necessary to define the correspondence between two surfaces. These examples shall demonstrate the algorithm at relatively simple data sets since in case of more complex input surfaces it is often hard to imagine how the Minkowski sum actually looks like.

Let A be a piecewise smooth surface which contains edges and let B be a smooth surface. We restrict our consideration to a single edge c of A . Assume that c is a smooth curve which possesses the parametrization $\mathbf{c}(t)$. The convolution of the curve c and the surface B is in fact part of the envelope which is obtained by moving B along c . This one-parameter family of motions consist of pure translations only. The instantaneous velocity vector field is determined by the derivative $\dot{\mathbf{c}}(t)$ of $\mathbf{c}(t)$ with respect to t . Let $\mathbf{n}_l(t)$ and $\mathbf{n}_r(t)$ be the left and right hand side unit normals of A at the curve c . These two vector fields define a wedge which follows the curve c . The contribution of the edge c and surface B to the outer boundary of the convolution consists of those part of $\mathbf{c} + B$ which lies inside the wedge spanned by vector fields $\mathbf{n}_l(t), \mathbf{n}_r(t)$.

4.3 Global self intersections of the envelope

So far we have computed the convolution surface $C = A \star B$ of two input surfaces A, B and we have taken care of edges and corners. The outer boundary Γ of C has been extracted

and is covered by data points $\mathbf{g} \in \Gamma$ and associated unit normals \mathbf{n}_g . Further we assume that T is a valid triangulation of the point set Γ .

According to the nature of envelopes, C and also Γ might possess global self intersections and we present a method to detect and to model these features. The algorithm for detection of sharp features on Γ is motivated by the following observations:

- The smooth parts of the envelope Γ possesses a relatively uniform sampling by the point set Γ since data points \mathbf{g} are obtained by searching for nearest neighbors with respect to a set of grid points of the stopping active front, see Section 4.1.
- The distances of points \mathbf{g} in smooth parts of Γ do not vary much and thus the size of triangles and the length of their segments is nearly constant.

According to these observations we detect sharp features by investigating distances of adjacent points $\mathbf{g}_1, \mathbf{g}_2$ taking into account the variation of their unit normal vectors $\mathbf{n}_1, \mathbf{n}_2$. For simplicity we have ignored the subscript g for the unit normals.

Let $\mathbf{g} \in \Gamma$ and \mathbf{n} an arbitrary point and its (unit) normal vector. We form the set $\hat{\Gamma} \subset \mathbb{R}^6$ consisting of points $\hat{\mathbf{g}} = (\mathbf{g}, w\mathbf{n})$, where w is a suitable chosen constant scaling the unit normals according to the average length of segments within the triangulation.

The canonical metric in \mathbb{R}^6 induces a metric on $\hat{\Gamma}$ and this defines a *feature sensitive* metric on Γ (and thus on Γ). See [13] for more information concerning the feature sensitive metric.

Once having defined the feature locations on Γ we form connected components (on the triangulation T) and investigate the neighborhood of each connected component (with respect to the triangulation) of the sharp features. Each connected component is replaced by a spine curve which represents the shape of the sharp feature.

Then profile sections of the investigated neighborhoods, orthogonal to the spine curve are computed. For typical self intersections we can recognize different sides of the neighborhood of a sharp feature. The intersection points of the approximating lines of the profile sections are good estimates for the true self intersection. The intersection points are sorted and smoothed by a spline curve which serves as approximation of the sharp feature. After having inserted a regularly sampled sequence of points of each of these spline curves into the data set Γ , a new triangulation is performed, see Figure 8.

4.4 Remarks on the implementation

The test implementation has been performed in Matlab and uses triangulated data as input. The models presented here are described by 2.5k to 10k data points. Additionally to points and triangles we store a list of edges and adjacency information for points as well as for triangles. Since the models are densely sampled, the correspondence between points

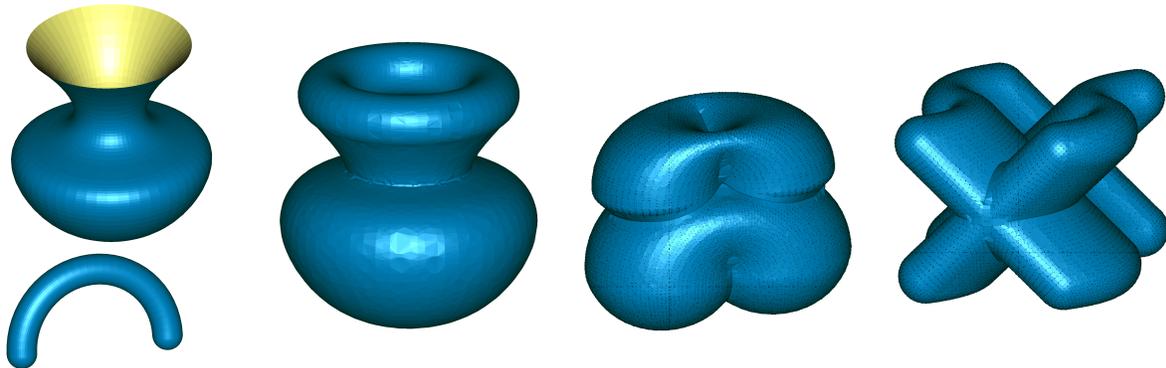


Figure 8: Cad-models, offset of the vase, Minkowski sum of pipe surface and vase, and object with sharpened edges from Figure 7 .

\mathbf{a} , \mathbf{b} of two models has been computed by searching nearest neighbors between the sets of unit normal vectors \mathbf{n}_a , \mathbf{n}_b .

The grid used for evaluation of the unsigned distance function has a size of δ , the average point distance on the convolution surface. We use a regular grid within the bounding box of the envelope. According to the size of the memory we are using, the number of grid points is limited by ~ 1.5 m.

The computation takes usually less than 60 sec., including nearest neighbor searching, initialization and marching on the grid, extraction of outer boundary, etc., if the data structure for the models is already prepared. The program is a prototypical implementation and thus it is not optimized concerning run-time and memory management.

Acknowledgment. This work has been funded by the Austrian Science Fund FWF within the research network S92.

References

- [1] Bajaj, C. and Kim, M.-S. 1989. Generation of configuration space obstacles: The case of a moving algebraic curve, *Algorithmica* 4(2), 157–172.
- [2] Farin, G., Hoschek, J., and Kim, M.-S., 2002. *Handbook of Computer Aided Geometric Design*, Elsevier.
- [3] Farouki, R.T., 2003. Minkowski combination of complex sets - Geometry, Algorithms and applications, in *Curve and Surface Design: Saint Malo 2002*, T. Lyche, M.-L. Mazure and L.L. Schumaker (eds.), Nashboro Press, 123–146.

- [4] Halperin, D., 2002. Robust Geometric Computing in Motion, The international Journal of Robotics Research, Vol. 21, No. 3, pp.219–232.
- [5] Kaul, A. and Farouki, R.T., 1995. Computing Minkowski sums of planar curves, *Int. J. of Computat. Geom. and Appl.* **5**, 413–432.
- [6] Kohler, K. and Spreng, M., 1995. Fast Computation of the C-Space of Convex 2D Algebraic Objects, *The Int. J. of Robotics Research* **14**(6), 590-608.
- [7] Lee, I.-K., Kim, M.-S. and Elber, G., 1998a. Polynomial/Rational Approximation of Minkowski Sum Boundary Curves, *Graphical Models* **60**, No.2, 136–165.
- [8] Lee, I.K., Kim, M.-S. and Elber, G., 1998b. The Minkowski Sum of 2D Curved Objects, *Proceedings of Israel-Korea Bi-National Conference on New Themes in Computerized Geometrical Modeling*, Tel-Aviv University, 155-164.
- [9] Mühlthaler, H. and Pottmann, H., 2003. Computing the Minkowski sum of ruled surfaces, *Graphical Models*, **65**, 369–384.
- [10] O’Rourke, J., 1994: *Computational Geometry in C*, Cambridge University Press.
- [11] Peternell, M. and Manhart, F., 2003. The convolution of a paraboloid and a parametrized surface, *Journal for Geometry and Graphics* **7**, 157-171.
- [12] Peternell, M., Pottmann, H., Steiner, T. and Zhao, H., 2005: Swept Volumes, *Computer-Aided Design Appl.* **2**, pp. 599-608.
- [13] Pottmann, H., Steiner, T., Hofer, M., Haider, C. and Hanbury, A., 2004. The isophotic metric and its application to feature sensitive morphology on surfaces, in: *Computer Vision — ECCV 2004, Part IV*, eds.: Pajdla, T. and Matas, J., Lecture Notes in Computer Science, Springer, **3024**, pp. 560–572.
- [14] Sampoli, M.L., Peternell, M. and Jüttler, B., 2005. The convolution of a paraboloid and a parametrized surface, *Computer Aided Geometric Design* **xx**, to appear.
- [15] Seong, J.-K., Kim, M.-S. and Sugihara, K., 2002. The Minkowski Sum of Two Simple Surfaces Generated by Slope-Monotone Closed Curves, *Proceedings of Geometric Modeling and Processing 2002*, Japan, 33–42.
- [16] Varadhan, G. and Manocha, D., 2004. Accurate Minkowski Sum Approximation of Polyhedral Models, In *Pacific Conference on Computer Graphics and Applications*, IEEE Computer Society, pp. 392–401.
- [17] Kim, Y.J., Varadhan, G., Lin, M.C. and Manocha, D., 2004. Fast swept volume approximation of complex polyhedral models, *Computer Aided Design* **36**, pp. 1013–1027. IEEE Computer Society, pp. 392–401.

- [18] Wallner, J. and Yang, Q., 2005. Swept volumes of many poses, Geometry Processing 2005, Symposium Proceedings, Eurographics/ACM SIGGRAPH, pp.161–170.