

Deformable Smooth Surface Design

H. Edelsbrunner

Department of Computer Science, University of Illinois,
Urbana, IL 61801, USA
edels@cs.uiuc.edu
and
Raindrop Geomagic Inc.
Champaign, IL 61820, USA

Abstract. A new paradigm for designing smooth surfaces is described. A finite set of points with weights specifies a closed surface in space referred to as *skin*. It consists of one or more components, each tangent continuous and free of self-intersections and intersections with other components. The skin varies continuously with the weights and locations of the points, and the variation includes the possibility of a topology change facilitated by the violation of tangent continuity at a single point in space and time. Applications of the skin to molecular modeling and to geometric deformation are discussed.

1. Introduction

This paper introduces a new paradigm for the design of smooth surfaces. The method works in all fixed dimensions, d , and generates closed $(d - 1)$ -manifolds in \mathbb{R}^d . For pragmatic reasons the paradigm is described in $d = 3$ dimensions and most figures illustrate the case $d = 2$.

Ideas and Concepts. The paradigm constructs surfaces as zero-sets of differentiable maps. This is an old idea expressed in the preimage theorem of differential topology, see, e.g., p. 21 of [11]. The crucial questions are computational: How do we specify such a map compactly and how do we construct the zero-set efficiently? We describe a mechanism that derives a differential map $f: \mathbb{R}^3 \rightarrow \mathbb{R}$ from a finite set of weighted points in \mathbb{R}^3 .

First we discuss how the zero-set, $f^{-1}(0)$, can be constructed directly from the finite specification, without computing f . An algebra is described that generates an infinite family of spheres from a finite set of weighted points with the property that the envelope




Fig. 1. The skin surface defined by nine weighted points in \mathbb{R}^3 , eight at the corners and one at the center of a cube.

of the spheres is the zero-set. We call this envelope the *skin* and the subset of \mathbb{R}^3 bounded by the envelope the *body* of the set of weighted points, see Fig. 1. Second we consider the construction of f and the deformation of the skin by taking a continuous sequence of preimages $f^{-1}(\tau)$. The flexibility in choosing different maps f with the same zero-set translates into the freedom of deforming the skin in different ways. The topology of the skin changes when τ passes a critical value.

Summary of Results. Whether or not skin and body are indeed useful concepts in geometric modeling depends on the existence of efficient algorithms. We describe a discrete framework and combinatorial algorithms constructing the skin as a collection of quadratic patches. The framework is based on Voronoi, Delaunay, and Alpha complexes of a finite set of points with weights [7]. The geometric and computational properties of skin and body are summarized in the following nonexhaustive list of informal claims:

- S1. decomposability: skin consists of a finite number of degree-2 patches (Theorem 13, Section 6),
- S2. constructibility: there are fast combinatorial algorithms constructing skin (Section 6),
- S3. symmetry: skin can be defined from the inside as well as the outside (Theorem 14, Section 6),
- S4. smoothness: in the nondegenerate case skin is everywhere tangent continuous (Theorem 16, Section 7),
- S5. economy: even a small number of weighted points can generate fairly complicated skin,
- S6. universality: every orientable closed surface has a skin representation,
- S7. deformability: topology changes of skin can be efficiently computed (Theorem 18, Section 7),
- S8. continuity: skin varies continuously with points and weights (Section 8).

An important project for the future is the comparison of the skin paradigm with other surface design methods described in the literature, see [10].

Applications. Surface design is a large part of computer-aided design and modeling, CAD/CAM, which is an industry that reaches into many aspects of modern life. Think of the design of automobile, aircraft, or ship parts or of geometric models for computer-animated movies, etc. This paper describes a new paradigm for surface design and is therefore relevant to all these activities. The question is whether the paradigm can compete with the currently prevalent spline methods, see [10] for a comprehensive introduction and [18] for one of the more recent extensions. The difference between the skin paradigm and conventional spline methods relates to the global versus local dichotomy. The skin paradigm defines a surface globally as the zero-set of a differential map; splines are implicit and derived automatically from the global definition. Conventional spline methods form a surface by gluing together locally parametrized patches. An important task now is to identify specific applications for which the skin paradigm offers sufficiently many advantages to warrant a large-scale departure from traditional methods.

An application where the superior properties of skin outweigh our hesitation to venture a premature prediction is the study of small and large molecules. We argue that skin should be substituted for the similar but geometrically ambiguous molecular surface model [4], [19]. In this context we mention the symmetry of skin with respect to inside and outside. This property begs to be exploited in problems on shape complementarity, such as molecular docking which is important in drug design, see, e.g., [12]. Finally, the continuous deformation of skin can be used to connect snapshots of molecular dynamics simulations and to support rational arguments about change and correspondence over time.

Outline. Section 2 introduces the algebra of spheres. Section 3 applies the algebra to create and study infinite sphere families generated from just two or three weighted points. Section 4 defines skin and body for arbitrary finite sets of weighted points. Section 5 introduces Voronoi and Delaunay complexes. Section 6 shows how skin can be decomposed into finitely many patches of degree 2. Section 7 proves topological properties of skin and body. Section 8 addresses the question of deformation and the equivalent problem of map design. Section 9 discusses applications of skin and body to molecular modeling. Section 10 mentions directions for further research.

2. Algebra of Spheres

Skin and body have been defined in an earlier publication by the author [8]. The algebraic structures that we are about to introduce greatly simplify the description and warrant a complete revision of [8], which is this paper. The origins of the algebraic structures can be traced back to independent work by Clifford, Darboux, and Frobenius after the middle of the last century, see the survey article [17]. The recent text by Pedoe [16] offers a detailed treatment of the algebra of circles and many applications.

Spheres Lifted to Points. The norm of a point $x = (\xi_1, \xi_2, \xi_3)$ in \mathbb{R}^3 is

$$\|x\| = \left(\sum_{i=1}^3 \xi_i^2 \right)^{1/2}.$$

We describe a *weighted point* $p \in \mathbb{R}^3 \times \mathbb{R}$ by its *location* $z_p \in \mathbb{R}^3$ and its *weight* $w_p \in \mathbb{R}$. The (*weighted*) *distance* from p is a map $\pi_p: \mathbb{R}^3 \rightarrow \mathbb{R}$ defined by

$$\pi_p(x) = \|x - z_p\|^2 - w_p.$$

For zero weight this is the square of the Euclidean distance between x and z_p . The zero-set of π_p is the sphere with center z_p and radius $\sqrt{w_p}$, and it is convenient to make no distinction between the weighted point and this sphere. For negative weight we get an imaginary sphere with imaginary radius. Imaginary spheres are of paramount importance in our theory of shape and surface. They form the key to mastering nonconvexity by filling and lending structure to the complement space. This becomes clear and rather concrete in later sections of this paper.

We base our algebraic structures on a bijection $\Pi: \mathbb{R}^3 \times \mathbb{R} \rightarrow \mathbb{R}^4$ that maps each sphere in \mathbb{R}^3 to a point in \mathbb{R}^4 . For $z_p = (\zeta_1, \zeta_2, \zeta_3)$ we have

$$\Pi(p) = (\zeta_1, \zeta_2, \zeta_3, \|z_p\|^2 - w_p).$$

For example, if p has radius $r_p = w_p^{1/2} = 0$, then $\Pi(p)$ lies on the paraboloid $\varpi: \mathbb{R}^3 \rightarrow \mathbb{R}$ defined by

$$\varpi(x) = \|x\|^2.$$

If p has positive real radius, then $\Pi(p)$ lies below ϖ and if it has imaginary radius, then $\Pi(p)$ lies above ϖ .

Vector Space. \mathbb{R}^4 together with the usual componentwise addition and scalar multiplication forms a vector space. We use the bijection Π to pull this vector space down to the space of spheres.

Let p and q be spheres in \mathbb{R}^3 and let $\gamma \in \mathbb{R}$ be a scalar. We define spheres $p + q$ and $\gamma \cdot p$ so that

$$\begin{aligned} \Pi(p + q) &= \Pi(p) + \Pi(q), \\ \Pi(\gamma \cdot p) &= \gamma \cdot \Pi(p), \end{aligned}$$

where the operations on the right are the usual vector space operations for points in \mathbb{R}^4 . The projection of the operations to \mathbb{R}^3 , the space of centers, is the usual vector space of \mathbb{R}^3 with $z_{p+q} = z_p + z_q$ and $z_{\gamma \cdot p} = \gamma \cdot z_p$. The projection to \mathbb{R} , the space of weights, is more complicated. Let $\langle \cdot, \cdot \rangle: \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}$ be the usual scalar product.

Fact 1. *The weights of $p + q$ and $\gamma \cdot p$ are*

$$\begin{aligned} w_{p+q} &= w_p + w_q + 2\langle z_p, z_q \rangle, \\ w_{\gamma \cdot p} &= \gamma w_p + (\gamma^2 - \gamma)\|z_p\|^2. \end{aligned}$$

The two formulas can be verified by straightforward algebraic manipulations. We call the set of spheres together with addition and scalar multiplication as defined the *vector space of spheres*.

Changing Size. Addition and scalar multiplication can be used to shrink or grow a sphere without changing its center. Let p be a sphere and let p' be the sphere with the same center, $z_{p'} = z_p$, and with zero weight. For a parameter $s \in \mathbb{R}$, define

$$p^s = (1 - s) \cdot p' + s \cdot p.$$

Use Fact 1 to determine the centers and weights of $p_1 = (1 - s) \cdot p'$ and $p_2 = s \cdot p$:

$$\begin{aligned} z_{p_1} &= (1 - s) \cdot z_p, \\ z_{p_2} &= s \cdot z_p, \\ w_{p_1} &= (s^2 - s) \|z_p\|^2, \\ w_{p_2} &= s w_p + (s^2 - s) \|z_p\|^2. \end{aligned}$$

Again use Fact 1 to add p_1 and p_2 :

Fact 2. *The center and weight of p^s are $z_{p^s} = z_p$ and $w_{p^s} = s \cdot w_p$.*

The radius of p^s is \sqrt{s} times the radius r_p of p . If r_p is a positive real, then $1 < s$ means growth, $0 \leq s < 1$ means shrinking, and $s < 0$ makes the radius imaginary. If r_p is imaginary the effect of the operation is symmetric.

Distance and Orthogonality. We extend the notion of weighted distance to pairs of weighted points or spheres:

$$\pi_{p,q} = \|z_p - z_q\|^2 - w_p - w_q.$$

Spheres p and q are *orthogonal*, denoted $p \perp q$, if $\pi_{p,q} = 0$. Indeed, for spheres with positive radius this is the condition that the two spheres meet in a circle and the two tangent planes passing through a common point x on this circle form a right angle. p and q are *further than orthogonal* if $\pi_{p,q} > 0$. For each sphere p let $\eta_p: \mathbb{R}^3 \rightarrow \mathbb{R}$ be defined by

$$\begin{aligned} \eta_p(x) &= \|x\|^2 - \pi_p(x) \\ &= 2\langle x, z_p \rangle - \|z_p\|^2 + w_p, \end{aligned}$$

whose graph is a hyperplane in \mathbb{R}^4 . If $w_p = 0$, then η_p is tangent to the paraboloid ϖ . More generally, the sphere p in \mathbb{R}^3 is the projection of the intersection between the paraboloid ϖ and the hyperplane η_p , see Fig. 2. The points of ϖ below and above η_p project inside and outside the sphere p . The algebraic expression of this observation is a trivial rearrangement of terms in the definition of η_p .

Fact 3. *Let p be a sphere and let x be a point in \mathbb{R}^3 . Then $\pi_p(x) = \varpi(x) - \eta_p(x)$.*




Fig. 2. A circle p in \mathbb{R}^2 corresponds to a plane η_p in \mathbb{R}^3 so p is the projection of $w \cap \eta_p$. The plane is not explicitly shown. Furthermore, $w \cap \eta_p$ bounds the patch on w visible from $\Pi(p)$ below w , and the patch projects to the disk bounded by p .

Orthogonality between spheres translates to incidence between points and hyperplanes:

Fact 4. $p \perp q$ iff $\Pi(p) \in \eta_q$ iff $\Pi(q) \in \eta_p$.

Consider, for example, the case where $w_p > 0$ and $w_q = 0$. Then $p \perp q$ iff z_q lies on the sphere p . The hyperplane η_q is tangent to w and touches the paraboloid in point $\Pi(q)$. By Fact 4, η_q passes through point $\Pi(p)$. In geometrical terms this implies that $w \cap \eta_p$ is the silhouette of the patch on w visible from $\Pi(p)$ below w , see Fig. 2.

Shrinking Orthogonal Spheres. Consider shrinking two spheres using scalars $s, t \geq 0$ with $s+t = 1$. Orthogonality is the limiting case in which the shrinking pulls intersecting spheres apart, no matter what their initial radii are, see Fig. 3.

Lemma 5. Let $p \perp q$ with real radii r_p and r_q and let $s, t \geq 0$ with $s + t = 1$. Then

$$p^s \cap q^t = \begin{cases} x & \text{if } sr_q^2 = tr_p^2, \\ \emptyset & \text{otherwise,} \end{cases}$$

where $x = (r_q^2 \cdot z_p + r_p^2 \cdot z_q) / (r_p^2 + r_q^2)$.




Fig. 3. Orthogonal circles p and q with equal radii $r_p = r_q$. Circles $p^{1/2}$ and $q^{1/2}$ touch at the midpoint between z_p and z_q . The dashed circles, $p^{1/4}$ and $q^{3/4}$, are disjoint.

Proof. Since $p \perp q$ and $s + t = 1$ we have

$$\begin{aligned} \|z_p - z_q\|^2 &= r_p^2 + r_q^2 \\ &\geq (s + t)(r_p^2 + r_q^2) - (\sqrt{s}r_q - \sqrt{t}r_p)^2 \\ &= (\sqrt{s}r_p + \sqrt{t}r_q)^2. \end{aligned}$$

Taking roots left and right implies that the radii of p^s and q^t add up to at most the Euclidean distance between the centers. We have equality iff $\sqrt{s}r_q = \sqrt{t}r_p$ or equivalently $sr_q^2 = tr_p^2$, as claimed. In this case

$$s = \frac{r_p^2}{r_p^2 + r_q^2} \quad \text{and} \quad t = \frac{r_q^2}{r_p^2 + r_q^2},$$

and the two shrunken spheres meet at the point $x = t \cdot z_p + s \cdot z_q$ specified in the claim. \square

3. Flats of Spheres

The flat defined by a set of points in \mathbb{R}^4 is the lowest-dimensional affine subspace that contains the set. This section extends this concept to spheres in \mathbb{R}^3 .

Affine Combination. The *affine hull* of a set of spheres $P = \{p_0, p_1, \dots, p_n\}$ is

$$\text{aff } P = \left\{ p = \sum_{i=0}^n \gamma_i \cdot p_i \mid \sum_{i=0}^n \gamma_i = 1 \right\},$$

see Fig. 4. Each sphere $p \in \text{aff } P$ is an *affine combination* of the p_i . P is *affinely independent* if $p_i \notin \text{aff}(P - \{p_i\})$ for every i . The maximum cardinality of any affinely independent set of spheres in \mathbb{R}^3 is 5. A *k-flat* is the affine hull of an affinely independent set P of cardinality $k + 1$. Its *dimension* is $\dim \text{aff } P = k$. Flats of dimension 1 or higher are simple examples of infinite sphere families.

For a set of spheres F and a parameter $s \in \mathbb{R}$ define

$$F^s = \{p^s \mid p \in F\}.$$

If $F = \text{aff } P$ and $\text{card } P \geq 2$, then F^s is an infinite family of spheres. Observe that shrinking and taking the affine hull does not commute. We are interested in the envelope of an infinite family of shrunken spheres. To convert spheres into balls let

$$\text{ucl } p = \{q \mid z_q = z_p, r_q^2 \leq r_p^2\}$$

and define the *upward closure* of a set of spheres F :

$$\text{ucl } F = \bigcup_{p \in F} \text{ucl } p.$$




Fig. 4. Two orthogonal 1-flats of circles in \mathbb{R}^2 . Every circle in the first flat is orthogonal to every circle in the second flat.

Keep in mind that $\text{ucl } F$ is a *set* of spheres; the *union* of these spheres is the same as the union of balls bounded by spheres in F . The *envelope* of F^s is the boundary of the union of balls:

$$\text{env } F^s = \text{bd } \bigcup \text{ucl } F^s,$$

see Fig. 1.

Orthogonal Flats. The set of spheres orthogonal to a single sphere p is $p^\perp = \Pi^{-1}(\eta_p)$, see Fact 4. η_p is a three-dimensional affine subspace of \mathbb{R}^4 , and accordingly p^\perp is a 3-flat of spheres in \mathbb{R}^3 . Indeed, for every point $z_q \in \mathbb{R}^3$ there is a unique radius r_q so that p and q are orthogonal. More generally, the *orthogonal flat* of a set of spheres P is

$$\begin{aligned} P^\perp &= \{q \mid q \perp p \text{ for all } p \in P\} \\ &= \bigcap_{p \in P} p^\perp. \end{aligned}$$

If q is orthogonal to all $p \in P$, then it is orthogonal to all $p \in \text{aff } P$ and therefore $P^\perp = (\text{aff } P)^\perp$. By the symmetry of Fact 4, a point p is orthogonal to every $q \in P^\perp$ iff $p \in \text{aff } P$. In short, $\text{aff } P = (P^\perp)^\perp$. The dimension of P^\perp is $\dim P^\perp = 3 - \dim \text{aff } P$. In \mathbb{R}^2 , the dimensions of orthogonal flats add up to 2, see Fig. 4.

Let $F = \text{aff } P$ and $G = P^\perp$ be two orthogonal flats and consider the sets of sphere centers:

$$\begin{aligned} F' &= \{z_p \mid p \in F\}, \\ G' &= \{z_q \mid q \in G\}. \end{aligned}$$

In the assumed case that $P' = \{z_p \mid p \in P\}$ is affinely independent, F' and G' are orthogonal affine subspaces of \mathbb{R}^3 . The dimensions add to

$$\begin{aligned} \dim F' + \dim G' &= \dim F + \dim G \\ &= 3. \end{aligned}$$

It follows that the two subspaces meet at a point $x(F) = F' \cap G'$ called the *focus* of F . Because orthogonality is symmetric, the focus of F is also the focus of G : $x(F) = x(G)$. For example, in Fig. 4 the focus is the intersection point of the two orthogonal lines that carry the centers of the circles. Let $p_x \in F$ and $q_x \in G$ be the two spheres with center $x = x(F)$. We have $p_x \perp q_x$ and therefore

$$r_{p_x}^2 + r_{q_x}^2 = 0.$$

So p_x is a real sphere iff q_x is imaginary. Similarly, p_x is a degenerate sphere or point iff q_x is a point. The orthogonal flats F and G are completely determined by F' , G' , and the radius of p_x . Note that $z_{p_x} = x$ is closest to $z_{q_x} = x$ among all points in F' and recall that all spheres in F are orthogonal to q_x . It follows that p_x has the smallest square radius of any sphere in F . Symmetrically, q_x has the smallest square radius of any sphere in G .

Complementarity. We consider shrunken versions of the orthogonal flats F and G :

$$\begin{aligned} F^s &= \{p^s \mid p \in F\}, \\ G^t &= \{q^t \mid q \in G\}. \end{aligned}$$

If s and t satisfy the requirements of Lemma 5, then F^s and G^t share the same envelope and together they cover the entire space:

Lemma 6. *Let $s, t > 0$ with $s + t = 1$. Then*

$$\begin{aligned} \bigcup \text{ucl } F^s \cup \bigcup \text{ucl } G^t &= \mathbb{R}^3, \\ \bigcup \text{ucl } F^s \cap \bigcup \text{ucl } G^t &= \text{env } F^s \\ &= \text{env } G^t. \end{aligned}$$

Proof. Let F' and G' be the affine subspaces of centers, let $x = F' \cap G'$, and consider $p_x \in F$ and $q_x \in G$. Recall that $r_{p_x}^2 + r_{q_x}^2 = 0$ and assume $r_{p_x}^2 \geq 0$. Since p_x has minimum square radius in F , every $p \in F$ has nonnegative square radius. The spheres $q \in G$ for which q^t touches p^s satisfy

$$r_q^2 = \frac{t}{s} \cdot r_p^2,$$




Fig. 5. The circles in Figure 4 after shrinking with $s = t = \frac{1}{2}$. The envelope is a hyperbola that separates the two families of circles.

which is again nonnegative, see Lemma 5. These spheres form a subset of the $(3 - k)$ -flat G , and the geometry of this subset is a $(2 - k)$ -dimensional sphere around x . The shrunken versions of these spheres touch p^s along a $(2 - k)$ -sphere in \mathbb{R}^3 . The union of the $(2 - k)$ -spheres, over all $p^s \in F^s$, is a manifold of dimension $(2 - k) + k = 2$. This manifold or surface is the common envelope of F^s and G^t . The spheres of F^s and G^t touch the surface from opposite sides, and each family covers the entire space on its side of the surface. \square

A side benefit of the above argument is the observation that a sphere q^t can touch the envelope only if its radius is sufficiently large. In particular, the smallest sphere q with q^t touching a sphere $p^s \in F^s$ has square radius

$$r_q^2 = \frac{t}{s} \cdot r_{p^s}^2,$$

which is positive if t and $r_{p^s}^2$ are positive.

Envelopes. We show that envelopes of flats are zero-sets of quadratic polynomials. First consider a 1-flat, F , and its orthogonal 2-flat, G . Suppose F' is the x_1 -axis and G' is the x_2x_3 -plane. Let r_0 be the radius of $p_0 \in F$ whose center is the focus of the two flats: $z_0 = 0 = F' \cap G'$. Let $s, t > 0$ with $s + t = 1$. We parametrize the family of spheres in F^s with τ and define

$$f(\tau, x_1, x_2, x_3) = (x_1 - \tau)^2 + x_2^2 + x_3^2 - s(\tau^2 + r_0^2).$$

The envelope consists of the points $x = (x_1, x_2, x_3)$ that satisfy $f = \partial f / \partial \tau = 0$. Observe that

$$\frac{\partial f}{\partial \tau}(\tau, x_1, x_2, x_3) = \tau(2 - 2s) - 2x_1.$$

From $\partial f / \partial \tau = 0$ we get $\tau = x_1 / (1 - s)$, and the envelope is the zero-set of

$$\frac{1}{s} f \left(\frac{x_1}{1-s}, x_1, x_2, x_3 \right) = -\frac{1}{t} x_1^2 + \frac{1}{s} (x_2^2 + x_3^2) - r_0^2.$$

Lemma 6 implies that this is also the envelope of G' . Alternatively, we can parametrize the spheres in G' and compute the envelope as before.

There are four types of envelopes in \mathbb{R}^3 distinguished by the dimension of F . Suppose F is a k -flat of spheres and F' is spanned by the first k coordinate axes. Then G is a $(3 - k)$ -flat and we can assume G' is spanned by the last $3 - k$ coordinate axes. Let r_0 be the radius of $p_0 \in F$.

Fact 7. *Let $s, t > 0$ with $s + t = 1$. The common envelope of F^s and G' is the zero-set of*

$$f(x_1, x_2, x_3) = -\frac{1}{t} \sum_{i=1}^k x_i^2 + \frac{1}{s} \sum_{j=k+1}^3 x_j^2 - r_0^2.$$

Indeed, for $k = 0$ the envelope is p_0^s itself, which is the zero-set of $f(x) = \|x\|^2 - sr_0^2$. For $k = 1$ the result agrees with the preceding calculation. The cases for $k = 2$ and $k = 3$ are symmetric to $k = 1$ and $k = 0$. The polynomial in Fact 7 implies a total of eight nondegenerate cases. They come in four pairs, each separated by a degenerate case, see Table 1.

Table 1. The different types of envelopes common to a shrunken k -flat and the shrunken orthogonal $(3 - k)$ -flat.*

k	$r_0^2 < 0$	$r_0^2 = 0$	$r_0^2 > 0$
0	Empty	Point	Sphere
1	2 sheets	Double-cone	1 sheet
2	1 sheet	Double-cone	2 sheets
3	Sphere	Point	Empty

*For $k = 1, 2$ we have hyperboloids of one and two sheets and a double-cone in the degenerate case. Each of the six different envelope types occurs twice with the spheres of the shrunken k -flat on opposite sides. In the case of the empty envelope this means that for $k = 0$ all spheres of the shrunken 0-flat are imaginary and for $k = 3$ the spheres of the shrunken 3-flat cover the entire space.

4. Skin and Body

This section extends the ideas explained for affine hulls in Section 3 to convex hulls. As it turns out the envelope of a shrunken convex hull of spheres consists of patches of envelopes of shrunken affine hulls. We begin with the necessary definitions.

Convex Combination. Recall that each sphere p in the affine hull of $P = \{p_0, p_1, \dots, p_n\}$ can be written as a sum of $\gamma_i \cdot p_i$, where the γ_i add up to 1. The *convex hull* of P is

$$\text{conv } P = \{p \in \text{aff } P \mid \gamma_i \geq 0 \text{ for all } i\}.$$

Figure 4 illustrates the convex hulls of two pairs of circles, one containing the leftmost and the rightmost circles, the other containing the topmost and the bottommost circles. Each $p \in \text{conv } P$ is a *convex combination* of the p_i . Clearly, $\text{conv } P \subseteq \text{aff } P$. Although the convex hull of P contains many more spheres than P , they are all contained in the union of balls defined by P :

Fact 8. $\bigcup \text{ucl } \text{conv } P = \bigcup \text{ucl } P.$

Let p be a convex combination of two spheres, p_0 and p_1 . To prove the claim it is sufficient to show that p is contained in the union of the two balls bounded by p_0 and p_1 . To see this is indeed the case observe that the sphere p passes through the circle $p_0 \cap p_1$ and that its center lies between the centers of p_0 and p_1 . If p_0 and p_1 are disjoint, then their common circle is imaginary and p is either imaginary or it is contained in one of the two balls.

The two most important new concepts in this paper are obtained from convex hulls by shrinking and taking the union and the envelope. For $0 \leq s \leq 1$ we define the *s-body* and the *s-skin* of P as

$$\begin{aligned} \text{bdy}^s P &= \bigcup \text{ucl } (\text{conv } P)^s, \\ \text{skn}^s P &= \text{bd } \text{bdy}^s P. \end{aligned}$$

See Fig. 1 for an example in \mathbb{R}^3 and Fig. 10 for an example in \mathbb{R}^2 .

Orthogonality and Complementarity. Recall that two spheres p and q are further than orthogonal from each other if $\pi_{p,q} > 0$. A sphere q that is orthogonal to or further than orthogonal from all spheres $p \in P$ has the same property with respect to all spheres p in the upward closure of the convex hull of P . The following definitions are therefore meaningful.

The set of spheres orthogonal to or further than orthogonal from a single sphere p is the *half-space* of spheres $\text{ucl } p^\perp$. We are interested in the intersection of all half-spaces defined by spheres $p \in P$. For finite P this is a convex polyhedron of spheres. The *orthogonal complement* of P is the smallest set of spheres Q , denoted as $\text{orth } P$, with

$$\text{ucl } \text{conv } Q = \bigcap_{p \in P} \text{ucl } p^\perp.$$

It contains points in \mathbb{R}^4 that correspond to spheres with arbitrarily large radii. To accom-




Fig. 6. P contains the four solid circles and $Q = \text{orth } P$ contains the two dashed circles and four dashed lines (infinitely large degenerate circles). The circles in Q have their centers at the Voronoi vertices of P .

modate these spheres, Q contains among others some infinitely large spheres or planes in \mathbb{R}^3 . In the nondegenerate case, each such plane is orthogonal to three spheres in P and further than orthogonal from all others. Similarly, each sphere with finite radius in Q is orthogonal to four spheres in P and further than orthogonal from all others, see Fig. 6. Observe that the construction of the orthogonal complement is symmetric. In other words, $P = \text{orth } Q = \text{orth orth } P$. Furthermore, the balls bounded by spheres in P and in Q cover the entire space:

Fact 9. $\bigcup_{p \in P} p \cup \bigcup_{q \in Q} q = \mathbb{R}^3$.

To see this suppose there is a point $x \in \mathbb{R}^3$ not covered by any of the balls. Then $(x, 0)$ is a sphere further than orthogonal to all $p \in P$ and not contained in the upward closure of $\text{conv } Q$, a contradiction to the definition of Q .

5. Proximity Complexes

The Delaunay complex [5] is the projection to the space of centers of the boundary complex of $\text{conv } P$. Similarly, the Voronoi complex [20] is the projection of the boundary complex of $\text{conv } Q$. This section introduces both complexes along with subcomplexes representing the shape bounded by the envelope of the spheres in P and in Q .

Voronoi Complex. Let P be a finite set of spheres in \mathbb{R}^3 and recall that $\pi_p(x) = \|x - z_p\|^2 - r_p^2$ is the weighted distance of a point x from a sphere p . The *Voronoi region* of $p \in P$ is

$$V_p = \{x \in \mathbb{R}^3 \mid \pi_p(x) \leq \pi_q(x) \text{ for all } q \in P\}.$$

In the assumed nondegenerate case V_p is either empty or a three-dimensional convex polyhedron. For a subset $X \subseteq P$ define

$$v_X = \bigcap_{p \in X} V_p.$$

By assumption of nondegeneracy, v_X is either empty or a convex polyhedron of dimension $\ell = \dim v_X = 4 - \text{card } X$. If v_X is nonempty we call it a *Voronoi ℓ -cell*. The *Voronoi*




Fig. 7. The Delaunay and Voronoi complexes of the four solid circles in Fig. 6. The Delaunay edges are solid and the Voronoi edges are dashed. Symmetrically, the dashed edges show the Delaunay complex and the solid edges show the Voronoi complex of the orthogonal complement.

complex is the collection of Voronoi cells:

$$\text{Vor } P = \{v_X \neq \emptyset \mid X \subseteq P\}.$$

The Voronoi complex decomposes the union of balls, $\bigcup \text{ucl } P$, into convex pieces. Indeed, if $x \in \bigcup \text{ucl } P$ belongs to V_p and is enclosed by a sphere $q \in P$, then $\pi_p(x) \leq \pi_q(x) \leq 0$ and x is also enclosed by p . This implies

$$V_p \cap \bigcup \text{ucl } P = V_p \cap \bigcup \text{ucl } p,$$

which is convex because V_p is convex and $\bigcup \text{ucl } p$ is the ball bounded by p and thus also convex, see Fig. 7.

Each Voronoi vertex or 0-cell $v = v_X$ has the same weighted distance $r^2 = \pi_p(v)$ from all four spheres $p \in X$. The sphere with center v and radius r is orthogonal to the $p \in X$, further than orthogonal from all other spheres in P , and it belongs to $Q = \text{orth } P$. In fact, Q contains a sphere for each Voronoi vertex and a plane (infinitely large degenerate sphere) for each unbounded Voronoi edge. The “center” of the plane lies at infinity in the direction of the corresponding Voronoi edge, see Fig. 7.

Delaunay Complex. For each Voronoi cell $v_X \in \text{Vor } P$ we have a *Delaunay cell*

$$\delta_X = \text{conv } X',$$

where $X' = \{z_p \mid p \in X\}$. The assumption of nondegeneracy implies that δ_X is a simplex in \mathbb{R}^3 . More precisely, δ_X is a Delaunay k -simplex iff v_X is a Voronoi ℓ -cell, where $k = 3 - \ell = \text{card } X - 1$. The *Delaunay complex* is the collection of Delaunay simplices:

$$\text{Del } P = \{\delta_X \mid v_X \in \text{Vor } P\}.$$

How does the Delaunay complex relate to P and to $Q = \text{orth } P$? Each Delaunay vertex or 0-simplex is the center of a sphere in P . The relation to Q is more intricate. Recall that each sphere $q \in Q$ with finite radius has its center at a Voronoi vertex v_X . $X \subseteq P$ contains four spheres and defines a Delaunay tetrahedron δ_X spanned by the centers of the spheres. The same Delaunay tetrahedron can be specified through inequalities:

$$\delta_X = \{y \in \mathbb{R}^3 \mid \pi_q(y) \leq \pi_p(y) \text{ for all } p \in Q\}.$$

In words, δ_X is the Voronoi region of $q \in Q$. This implies a fundamental symmetry between Voronoi and Delaunay complexes.




Fig. 8. The Delaunay subcomplex of the four solid circles in Fig. 6 consists of four vertices and two (solid) edges. The Voronoi subcomplex consists of two vertices and three (dashed) edges.

Fact 10. $\text{Del } P = \text{Vor } Q$ and $\text{Vor } P = \text{Del } Q$.

The statement is not entirely correct but we can make it correct by slightly altering the definition of the Delaunay complex. Specifically, we add the Voronoi cells defined by infinitely large spheres in Q to $\text{Del } P$. The resulting set of Delaunay cells decomposes the union of balls bounded by spheres in Q into convex pieces, see Fig. 7.

Subcomplexes. Recall that Voronoi regions decompose a union of balls into convex pieces. It follows that each Voronoi cell, v_X , forms a convex intersection with the union:

$$\begin{aligned} \varrho_X &= v_X \cap \bigcup \text{ucl } P \\ &= \bigcap_{p \in X} (V_p \cap \bigcup \text{ucl } p). \end{aligned}$$

Clearly $\varrho_X \subseteq v_X$, and by collecting all simplices that correspond to nonempty intersections we get a subcomplex of the Delaunay complex:

$$\text{Dsx } P = \{\delta_X \in \text{Del } P \mid \varrho_X \neq \emptyset\},$$

see Fig. 8. In the terminology of [7] and [9] $\text{Dsx } P$ is one of the alpha complexes of P , namely, the one defined for $\alpha = 0$.

Symmetrically, the simplices in $\text{Del } P$ decompose the union of balls bounded by spheres in the orthogonal complement. We define

$$\text{Vsx } P = \left\{ v_X \in \text{Vor } P \mid \delta_X \cap \bigcup \text{ucl } Q \neq \emptyset \right\},$$

which is a subcomplex of $\text{Vor } P$, see Fig. 8.

It can be shown that the union of balls bounded by spheres in P covers all simplices in $\text{Dsx } P$, see also Section 6. Symmetrically, the union of balls bounded by spheres in Q covers all cells in $\text{Vsx } P$. The two containment relations can be connected using Fact 9. To state the resulting chain of relations we denote the complement of a set $A \subseteq \mathbb{R}^3$ by $\overline{A} = \mathbb{R}^3 - A$.

Fact 11. *Assume P is in nondegenerate position. Then*

$$\bigcup \text{Dsx } P \subseteq \overline{\bigcup \text{ucl } Q} \subseteq \bigcup \text{ucl } P \subseteq \overline{\bigcup \text{Vsx } P}.$$

6. Decomposing Space and Skin

This section returns to the body and skin of a finite set of spheres in \mathbb{R}^3 . The main result is a decomposition of \mathbb{R}^3 into convex cells that decompose the skin into patches of algebraic degree 2.

Voronoi Regions in the Limit. We begin by revisiting k -flats of spheres. Let $P = \{p_0, p_1, \dots, p_k\}$ be an affinely independent set of spheres and assume that the set of centers, $P' = \{z_{p_0}, z_{p_1}, \dots, z_{p_k}\}$ is also affinely independent. Then $F = \text{aff } P$ is a k -flat of spheres and

$$\begin{aligned} F' &= \text{aff } P' \\ &= \{z_p \mid p \in F\} \end{aligned}$$

is a k -dimensional affine subspace of \mathbb{R}^3 . Let $s, t > 0$ be real parameters with $s + t = 1$, as usual. Let V_p^s be the Voronoi region of p^s among the infinitely many spheres in F^s . We construct V_p^s using a map $\varphi^s: F \rightarrow F'$. This map associates each sphere $p \in F$ with the point $y = \varphi^s(p) \in F'$ that satisfies

$$\pi_{p^s}(y) \leq \pi_{q^s}(y)$$

for all $q \in F$. V_p^s is the $(3 - k)$ -dimensional affine subspace orthogonal to F' with $F' \cap V_p^s = y$, see Fig. 9. The point $y = \varphi^s(p)$ lies between z_p and the focus $x = x(F)$ of F :

Fact 12. $\varphi^s(p) = s \cdot x + t \cdot z_p$.

In the limit when $s = 0$ we have $\varphi^s(p) = z_p$ for every $p \in P$. This corresponds to the case where all spheres in F^s are points and clearly every point is closest to itself. When s grows the points $\varphi^s(p)$ move closer to x and meet at x in the limit when $s = 1$. Fact 12 can be proved with analytical calculations similar to the ones preceding Fact 7 in Section 3.




Fig. 9. Skin construction for two circles in the plane. The vertical lines are Voronoi regions of circles in the convex hull reduced with $s = \frac{1}{2}$. The two extreme lines separate the circle pieces of the $\frac{1}{2}$ -skin from the connecting hyperbola piece.

Shrinking and Mixing. Fact 12 is significant because it can be used to decompose the skin defined by P . Before taking on the general case consider the two circles p_0 and p_1 in Fig. 9. The $\frac{1}{2}$ -skin has the shape of a dumbbell and consists of two circular arcs connected by a hyperbola piece between the lines $V_{p_0}^{1/2}$ and $V_{p_1}^{1/2}$.

We generalize the construction of the strip between the two lines. Let P be a finite set of spheres in \mathbb{R}^3 . For each $X \subseteq P$ with nonempty Voronoi cell v_X and corresponding Delaunay simplex δ_X define the *s-mixed cell*

$$\mu_X^s = s \cdot v_X + t \cdot \delta_X,$$

compare with Fact 12. For $0 < s < 1$ the dimension of μ_X^s is independent of the cardinality of X : $\dim \mu_X^s = \dim v_X + \dim \delta_X = 3$. There are four types of mixed cells: for card $X = 1, 2, 3, 4$ the mixed cell μ_X^s is a Voronoi 3-cell, a prism over a Voronoi 2-cell, a prism over a Delaunay triangle, and a Delaunay tetrahedron. Consider the collection of *s-mixed cells*:

$$\text{Mix}^s P = \{\mu_X^s \mid v_X \in \text{Vor } P\},$$

see Fig. 10. Observe that $\text{Mix}^0 P = \text{Del } P$, $\text{Mix}^1 P = \text{Vor } P$, and generally $\text{Mix}^s P = \text{Mix}^t Q$, where $Q = \text{orth } P$.

Here is an intuitive picture of the construction. Take the interval of three-dimensional affine subspaces of \mathbb{R}^4 defined by $x_4 = s$, for $s \in [0, 1]$. Draw $\text{Del } P$ in $x_4 = 0$ and $\text{Vor } P$ in $x_4 = 1$. For each Delaunay simplex and corresponding Voronoi cell construct

$$\mu_X = \text{conv}(\delta_X \cup v_X),$$

see Fig. 11. All μ_X are convex polyhedra of dimension 4, their interiors are mutually disjoint, and they decompose the strip between $x_4 = 0$ and $x_4 = 1$. The subspace $x_4 = s$ intersects μ_X in the *s-mixed cell*, μ_X^s .

Decomposition into Patches. Mixed cells are significant because they decompose the skin into patches of degree 2. These patches are pieces of envelopes of shrunken affine hulls. Recall that each envelope is the zero-set of a degree-2 polynomial, see Fact 7.




Fig. 10. The mixed cells for $s = \frac{1}{2}$ defined by the four circles in Fig. 6. There are three types: Voronoi regions, rectangles, and Delaunay triangles. The cells decompose the *s-skin* into circle and hyperbola pieces.




Fig. 11. The construction of mixed cells from the Delaunay complex in the bottom plane and the Voronoi complex in the top plane.

Theorem 13. For each $X \subseteq P$ and each $s \in [0, 1]$ we have

$$\mu_X^s \cap \text{skn}^s P = \mu_X^s \cap \text{env}(\text{aff } X)^s.$$

Proof. For each sphere $p \in \text{conv } P$ let V_p^s be the Voronoi region of p^s in the Voronoi complex of $(\text{conv } P)^s$. V_p^s is nonempty only if no other sphere with the same center has larger square radius than p . Assume p satisfies this maximum square radius criterion and let δ_X be the Delaunay simplex of lowest dimension that contains its center, z_p . Define $F = \text{aff } X$, which is a flat of spheres with dimension $k = \dim F$ and focus $x = x(F)$. Let $G = F^\perp$ be the orthogonal flat and let G' be the set of centers of spheres in G . G' is a $(3 - k)$ -dimensional affine subspace of \mathbb{R}^3 passing through the focus x of F . By Fact 12,

$$V_p^s \subseteq s \cdot G' + t \cdot z_p.$$

If $k = 3$, then G' consists of a single point, namely the focus x of F , and $V_p^s = s \cdot x + t \cdot z_p$ is the single point Voronoi region of p^s . The union of these points over all spheres p with $z_p \in \delta_X$ is exactly the s -mixed cell μ_X^s . In the other case, when $k < 3$, there are Delaunay simplices δ_Y that properly contain δ_X . These δ_Y constrain the relevant part of G' to within v_X , and we get $V_p^s = s \cdot v_X + t \cdot z_p$. To see this apply Fact 12 to the orthogonal complement and take the union over all spheres q with $z_q \in v_X$. Finally take the union of the reduced copies of v_X over all spheres p with $z_p \in \delta_X$. The result is again the s -mixed cell μ_X^s .

To summarize we showed that for every point $y \in \mu_X^s$ the spheres in $(\text{conv } P)^s$ that minimize the weighted distance from y belong to $(\text{conv } X)^s$. Similarly, the spheres in $(\text{aff } X)^s$ that minimize the weighted distance from y belong to $(\text{conv } X)^s$. In formulas:

$$\begin{aligned} \mu_X^s \cap \text{skn}^s P &= \mu_X^s \cap \text{env}(\text{conv } P)^s \\ &= \mu_X^s \cap \text{env}(\text{conv } X)^s \\ &= \mu_X^s \cap \text{env}(\text{aff } X)^s. \end{aligned} \quad \square$$

Theorem 13 is a technical statement of claim **S1** that skin is decomposable: the mixed cells decompose the skin into finitely many degree-2 patches. It also supports claim **S2** that skin is constructible: the mixed cells and patches are readily computed from the Voronoi complex and the Delaunay complex.

Complementarity. The s -mixed cells decompose the s -skin into finitely many patches. Within a cell μ_X^s only spheres defined by the Delaunay simplex δ_X and the Voronoi cell ν_X are relevant. In other words, within μ_X^s the s -skin of P looks the same as the envelope of $(\text{aff } X)^s$, and symmetrically, the t -skin of Q looks the same as the envelope of $(X^\perp)^t$. This is true for every mixed cell and together they cover the entire \mathbb{R}^3 . We conclude that the complementarity result stated in Lemma 6 for flats generalizes to convex hulls:

Theorem 14. *Let $s, t \geq 0$ with $s + t = 1$. Then*

$$\begin{aligned} \text{bdy}^s P \cup \text{bdy}^t Q &= \mathbb{R}^3, \\ \text{bdy}^s P \cap \text{bdy}^t Q &= \text{skn}^s P \\ &= \text{skn}^t Q. \end{aligned}$$

Theorem 14 is a technical statement of claim S3 that skin is symmetric: $\text{skn}^s P$ is defined from inside by the spheres in P and from outside by the spheres in $Q = \text{orth } P$.

Mixed Cell Classification. Some of the mixed cells are contained in the body of P , some intersect the skin, and the others lie outside the body. We derive information on this classification from the subcomplexes $\text{Dsx } P \subseteq \text{Del } P$ and $\text{Vsx } P \subseteq \text{Vor } P$ defined in Section 5.

The *boundary* of $\text{Dsx } P$ consists of all simplices contained in at least one Delaunay simplex not in $\text{Dsx } P$, and the *interior* consists of all other simplices:

$$\begin{aligned} \text{Bd Dsx } P &= \{\delta_X \in \text{Dsx } P \mid \delta_X \subseteq \delta_Y \notin \text{Dsx } P\}, \\ \text{Int Dsx } P &= \text{Dsx } P - \text{Bd Dsx } P. \end{aligned}$$

Similarly, the boundary and the interior of $\text{Vsx } P$ are $\text{Bd Vsx } P = \text{Bd Dsx } Q$ and $\text{Int Vsx } P = \text{Int Dsx } Q$. Note that δ_X belongs to the interior of $\text{Dsx } P$ iff ν_X is contained in $\bigcup \text{ucl } P$, and δ_X belongs to the boundary of $\text{Dsx } P$ iff ν_X has nonempty intersection with $\bigcup \text{ucl } P$ but is not contained in it.

Lemma 15. *Let P be a finite set of spheres in nondegenerate position in \mathbb{R}^3 .*

- (i) *If $\delta_X \in \text{Int Dsx } P$, then $\mu_X^s \subseteq \text{bdy}^s P$ for every $s \in [0, 1]$.*
- (ii) *If $\nu_X \in \text{Int Vsx } P$, then $\mu_X^s \cap \text{bdy}^s P = \emptyset$ for every $s \in [0, 1]$.*
- (iii) *If neither $\delta_X \in \text{Int Dsx } P$ nor $\nu_X \in \text{Int Vsx } P$, then $\mu_X^s \cap \text{skn}^s P \neq \emptyset$ for at least one $s \in [0, 1]$.*

Proof. Claims (i) and (ii) are symmetric and we prove only (i). We modify the construction of mixed cells illustrated in Fig. 11 so it handles intersections with spheres. We draw $\text{Dsx } P$ in $x_4 = 0$ and the convex cells $\varrho_X = \nu_X \cap \bigcup \text{ucl } P$ of Section 5 in $x_4 = 1$. By definition $\varrho_X \neq \emptyset$ iff $\delta_X \in \text{Dsx } P$ and we can take the convex hull:

$$\lambda_X = \text{conv}(\delta_X \cup \varrho_X).$$

Clearly, $\lambda_X \subseteq \mu_X$. If $\delta_X \in \text{Int Dsx } P$, then $\varrho_X = \nu_X$ and $\lambda_X = \mu_X$. Consider the cross section of the cells λ_X at $s \in [0, 1]$:

$$\lambda_X^s = s \cdot \varrho_X + t \cdot \delta_X.$$

The union of the cells λ_X^s contains the mixed cells μ_X^s that correspond to interior simplices δ_X . To complete the proof of (i) we show that the union is contained in the s -body of P . We rewrite the union of cells λ_X^s as a union of balls. Recall the definition of the focus $x = x(F)$ of $F = \text{aff } X$: it is the center of the sphere $p \in F$ with minimum square radius. Let r_X be the radius of p and observe that $r_X \geq 0$ for else δ_X would not be in $\text{Dsx } P$. For each point y in the interior of δ_X let b_y^s be the spherical ball with center y and radius $s \cdot r_X$. To simplify notation let $K = \bigcup \text{Dsx } P$. Then

$$\begin{aligned} \bigcup_{X \subseteq P} \lambda_X^1 &= \bigcup \text{ucl } P \\ &= \bigcup_{y \in K} b_y^1. \end{aligned}$$

The cross section at $s \in [0, 1]$ is

$$\begin{aligned} \bigcup_{X \subseteq P} \lambda_X^s &= \bigcup_{y \in K} b_y^s \\ &\subseteq \bigcup_{y \in K} b_y^{s\sqrt{s}} \\ &\subseteq \text{bdy}^s P, \end{aligned}$$

and (i) follows.

To prove (iii) we reformulate the premise: $\delta_X \not\subseteq \text{Int Dsx } P$ and $\nu_X \not\subseteq \text{Int Vsx } P$ is equivalent to $\nu_X \not\subseteq \bigcup \text{ucl } P$ and $\delta_X \not\subseteq \bigcup \text{ucl } Q$. Take points $y_1 \in \nu_X$ not in $\bigcup \text{ucl } P$ and $y_0 \in \delta_X$ not in $\bigcup \text{ucl } Q$. The union of skins $\text{skn}^s P$ over all $s \in [0, 1]$ forms an orientable 3-manifold in $\mathbb{R}^3 \times [0, 1]$. Points y_1 and y_0 lie on different sides of this manifold. Since y_1 and y_0 both belong to $\mu_X = \text{conv}(\delta_X \cup \nu_X)$ we conclude that μ_X intersects the manifold and there is an $s \in [0, 1]$ so that μ_X^s has nonempty intersection with $\text{skn}^s P$. \square

7. Topological Properties

The skin construction is unusual in the number and combination of mathematically elegant and practically useful properties. In this section we show that skin is tangent continuous, that different values of s define isotopic skins, and that the body is homotopy equivalent to the union of simplices in $\text{Dsx } P$.

Tangent Continuity. Recall that the cells in $\text{Mix}^s P$ decompose the s -skin into patches of the form

$$\varphi_X^s = \mu_X^s \cap \text{env}(\text{aff } X)^s,$$

see Theorem 13. We assume $s, t > 0$ with $s + t = 1$ as usual. The envelope that contains φ_X^s is the zero-set of a degree-2 function $f_X^s: \mathbb{R}^3 \rightarrow \mathbb{R}$. After suitable rotation we can write

$$f_X^s(x_1, x_2, x_3) = -\frac{1}{t} \sum_{i=1}^k x_i^2 + \frac{1}{s} \sum_{j=k+1}^3 x_j^2 - r_0^2,$$

see Fact 7. Except if $r_0 = 0$, the zero-set is everywhere tangent continuous. By this we mean that each point x in the zero-set of f_X^s defines a unique plane of tangential directions. This plane is normal to the gradient vector at x :

$$\nabla f_X^s(x) = \left(\frac{\partial f_X^s}{\partial x_1}(x), \frac{\partial f_X^s}{\partial x_2}(x), \frac{\partial f_X^s}{\partial x_3}(x) \right).$$

The plane is well defined unless $\nabla f_X^s(x)$ vanishes, which is the case only if x is the origin and $r_0 = 0$. We consider $r_0 = 0$ a degenerate case and for the time being assume it does not occur. In Section 8 this case resurfaces and plays an important role in deforming skin.

In the nondegenerate case each patch φ_X^s of the s -skin is tangent continuous in the interior. We give a geometric argument for this claim that also covers the boundary of φ_X^s . Define $f^s: \mathbb{R}^3 \rightarrow \mathbb{R}$ with $f^s(x) = f_X^s(x)$ for every $x \in \mu_X^s$ and recall that $\text{skn}^s P$ is the zero-set of f^s .

Theorem 16. *Let $0 < s < 1$ and assume P is a finite set of spheres in nondegenerate position in \mathbb{R}^3 . Then $\text{skn}^s P$ is tangent continuous.*

Proof. Consider a point x in the zero-set of f^s . By Theorem 14 there are spheres $p \in P$ and $q \in Q$ so p^s and q^t lie on opposite sides and touch the skin at x . By assumption of nondegeneracy both spheres have positive radius. The unique plane that separates p^s and q^t is the plane tangent to the zero-set at point x . \square

Theorem 16 is a more precise statement of claim S4 that skin is smooth: in the nondegenerate case skin is everywhere tangent continuous.

Skin Isotopy. Consider the interval of skins $\text{skn}^s P$ for $0 \leq s \leq 1$. Assume nondegeneracy so all skins except for $\text{skn}^0 P = \text{bd}(\bigcup \text{ucl } Q)$ and $\text{skn}^1 P = \text{bd}(\bigcup \text{ucl } P)$ are everywhere tangent continuous, see Fig. 12. The union of all skins is a channel:

$$\begin{aligned} H &= \bigcup_{0 \leq s \leq 1} \text{skn}^s P \\ &= \bigcup \text{ucl } P \cap \bigcup \text{ucl } Q. \end{aligned}$$




Fig. 12. Six skins in the interval $s \in [0, 1]$ sweeping out the channel between the boundaries of the disk unions defined by P and Q .

We form a fibration of H with fibers normal to the skins. Consider the middle skin, $\text{skn } P = \text{skn}^{1/2} P$, and for each point $x \in \text{skn } P$ define the fiber

$$g_x: [0, 1] \rightarrow H$$

with $g_x(\frac{1}{2}) = x$, $g_x(s) \in \text{skn}^s P$ for all $s \in [0, 1]$, and g_x normal to $\text{skn}^s P$ for all $s \in (0, 1)$. The fibers are the solutions or flow curves of the differential equation defined by a vector field $V: \text{int } H \rightarrow \mathbb{R}^3$. We have $V(x) = \nabla f^s(x)$ with $s \in (0, 1)$ chosen so x belongs to the zero-set of f^s . The vector field is continuous and does not vanish anywhere in the interior of H . We have a diffeomorphism

$$d^s: \text{skn } P \rightarrow \text{skn}^s P$$

defined by $d^s(x) = g_x(s)$ for every $s \in [0, 1]$.

The diffeomorphisms can be combined to form an isotopy between different skins. Assume without loss of generality that $\frac{1}{2} < s$. The isotopy between $\text{skn } P$ and $\text{skn}^s P$ consists of all intermediate diffeomorphisms:

$$i^s: \text{skn } P \times [\frac{1}{2}, s] \rightarrow \mathbb{R}^3,$$

with $i^s(x, r) = d^r(x)$. By transitivity we have an isotopy between every pair of skins.

Fact 17. *In the nondegenerate case there is an isotopy between $\text{skn}^r P$ and $\text{skn}^s P$ for every pair of parameters $0 < r, s < 1$.*

A possible application of the isotopy is the diffeomorphic maintenance of a pattern or drawing on the skin while the parameter s is altered.

Body Homotopy. The isotopy between skins can be used to construct isotopies between bodies. We prove a weaker result establishing that all bodies of P are homotopy equivalent. Given $s \in [0, 1]$ we simplify notation by defining $\mathbb{X} = \text{bdy}^0 P$ and $\mathbb{Y} = \text{bdy}^s P$. Recall that the body grows with the exponent and therefore $\mathbb{X} \subseteq \mathbb{Y}$. Let

$$D: \mathbb{Y} \times [0, 1] \rightarrow \mathbb{Y}$$

be defined so its restriction to \mathbb{X} is the identity for all $u \in [0, 1]$. For a point $y \in \mathbb{Y} - \mathbb{X}$ choose $r \in [0, 1]$ and $z \in \text{skn } P$ so that $y = g_z(r)$. Then

$$D(y, u) = g_z((1 - u)r),$$

for all $u \in [0, 1]$. D is a homotopy between the identity on \mathbb{Y} at $u = 0$ and a retraction from \mathbb{Y} to \mathbb{X} at $u = 1$. In words, D is a deformation retraction and \mathbb{X} and \mathbb{Y} are homotopy equivalent.

The homotopy result can be extended to the union of balls, $\bigcup \text{ucl } P$, and to the union of simplices in $\text{Dsx } P \subseteq \text{Del } P$ defined in Section 5. The homotopy equivalence between $\bigcup \text{Dsx } P$ and $\bigcup \text{ucl } P$ has been established in [7].

Theorem 18. *In the nondegenerate case all bodies $\text{bdy}^s P$, for $s \in [0, 1]$, are homotopy equivalent to $\bigcup \text{Dsx } P$.*

The significance of this result is primarily computational. $\text{Dsx } P$ is a combinatorial object amenable to fast algorithms. In particular, in \mathbb{R}^3 there is an algorithm that computes the betti numbers of $\bigcup \text{Dsx } P$ in time proportional to the cardinality of $\text{Dsx } P$ [6]. Theorem 18 implies that the betti numbers of $\bigcup \text{Dsx } P$ and those of the bodies $\text{bdy}^s P$ are the same.

Theorem 18 covers part of claim S7 that skin is deformable: the time, place, and way topology changes can be predicted by computing the change in the Delaunay subcomplex $\text{Dsx } P$.

8. Growth and Change

This section returns to the initial idea of constructing a differential map $f: \mathbb{R}^3 \rightarrow \mathbb{R}$ and using a continuous sequence of preimages to deform one shape into another. We begin with a deformation induced by growth.

Growing Spheres. For each sphere $p \in P$ and each parameter $\alpha^2 \in \mathbb{R}$ define $p_\alpha = (z_p, r_p^2 + \alpha^2)$. We allow α^2 to be negative in which case p_α is smaller than p and possibly imaginary. Define $P_\alpha = \{p_\alpha \mid p \in P\}$. Note that the weighted distance of a point $x \in \mathbb{R}^3$ is $\pi_p(x) - \alpha^2$. It follows that the Voronoi and Delaunay complexes do not vary with α : $\text{Vor } P_\alpha = \text{Vor } P$ and $\text{Del } P_\alpha = \text{Del } P$ for all $\alpha^2 \in \mathbb{R}$.

We fix $s \in (0, 1)$ and construct a differential map $f = f^s: \mathbb{R}^3 \rightarrow \mathbb{R}$ whose preimages $f^{-1}(\alpha^2)$ are the s -skins of the P_α . This is done by gluing patches of maps $f_X = f_X^s$ clipped to within their mixed cells μ_X^s . The result is the same map $f = f^s$ as in Section 7. By Fact 7, we can find a coordinate system so

$$f(x_1, x_2, x_3) = -\frac{1}{t} \sum_{i=1}^k x_i^2 + \frac{1}{s} \sum_{j=k+1}^3 x_j^2 - r_0^2$$

within μ_X^s . Since the growth model keeps the Voronoi and Delaunay complexes invariant, the map $F_\alpha(x) = f(x) - \alpha^2$ has zero-set $F_\alpha^{-1}(0) = \text{skn}^s P$. We can recover the s -skin of P_α from f directly by taking the preimage at α^2 :

Fact 19. *Let $\alpha^2 \in \mathbb{R}$ and $f: \mathbb{R}^3 \rightarrow \mathbb{R}$ as constructed. Then*

$$\begin{aligned} \text{skn}^s P_\alpha &= f^{-1}(\alpha^2), \\ \text{bdy}^s P_\alpha &= f^{-1}(-\infty, \alpha^2]. \end{aligned}$$

Changing Topology. To deform $\text{skn}^s P$ to $\text{skn}^s P_\beta$, for $\beta^2 > 0$, we generate a continuous interval of skins $f^{-1}(\alpha^2)$, for $\alpha^2 \in [0, \beta^2]$. It is possible that skin and body change topology as they go through the sequence. Specifically, this happens at critical points $y \in \mathbb{R}^3$ with vanishing gradient $\nabla f(y)$. These are precisely the focus points $x = x(\text{aff } X)$ that belong to the corresponding Delaunay simplex and Voronoi cell: $x = \delta_X \cap \nu_X$. In the nondegenerate case the intersection between δ_X and ν_X is either empty or a point in the interiors of both. In the latter case x is a critical point of f and it lies in the interior of μ_X^s .

Table 2. The topology of the body changes in analogy to adding a k -simplex to the corresponding Delaunay subcomplex, see Theorem 18.*

k	Body change	Skin change
0	Component appears	\emptyset to sphere
1	Bridge appears	2 sheets to 1 sheet
2	Tunnel disappears	1 sheet to 2 sheets
3	Void disappears	Sphere to \emptyset

*The change is described by switching from the first to the third column in Table 1. The second column in Table 1 describes the intermediate degenerate case.

The way the topology changes depends on the dimension $k = \dim \delta_X = 3 - \dim \nu_X$, see Table 2. Note that cases $k = 0$ and $k = 3$ are symmetric in terms of taking complements or switching from $\text{bdy}^s P$ to $\text{bdy}' Q$. Similarly, cases $k = 1$ and $k = 2$ are symmetric in the same sense. It follows that the topology of skin changes only in two ways:

- (i) A sphere appears or disappears.
- (ii) A hyperboloid of one sheet flips over to one of two sheets or vice versa.

Each topology change requires a momentary degenerate configuration of spheres. The change can be localized to a single point in time and space where the degeneracy leads to a violation of tangent continuity.

The type of a topology change is completely characterized by the sign sequence of the corresponding polynomial in Fact 7. Only the three signs of the quadratic terms are considered, and there are four triplets that differ even after permutations: $+++$, $-++$, $+-+$, $---$. There are $\dim \delta_X$ minus signs and $\dim \nu_X$ plus signs. The topology changes when the constant term of the polynomial, r_0^2 , passes through zero. Under the assumption of growth r_0^2 can only switch from negative to positive, not the other way round.

General Deformation. The growth-induced deformation can be generalized by allowing spheres to change size, move, and duplicate or merge all at the same time. A framework for such deformations is described in [2], and sample results in \mathbb{R}^2 and in \mathbb{R}^3 can be found in [3] and [2]. Rather than revisiting the details of this framework we briefly discuss the general idea from a Morse theoretic angle [14].

Consider a time interval $[0, 1]$, and for each time $\tau \in [0, 1]$ let P_τ be a set of spheres in \mathbb{R}^3 . We expect that P_τ satisfies necessary continuity requirements. Let $f: \mathbb{R}^3 \rightarrow \mathbb{R}$ be defined so that

$$\text{skn } P_\tau = f^{-1}(\tau)$$

for all $\tau \in [0, 1]$. The shape at time τ is the skin of P_τ , and it deforms as τ increases. The deformation can be studied with a standard Morse theoretical approach that requires f be everywhere differentiable and twice differentiable at critical points. It is also convenient

to assume all critical points are isolated, which is the nondegenerate case. Further details justifying claim S8 about continuous skin deformation are omitted.

9. Molecular Biology

Skin surfaces bear an unmistakable resemblance to space filling models of molecules used in computational biology literature. This section reviews these models and compares them with the skin as defined in this paper. We argue that skin has distinct advantages over the other models.

Space Filling Models. A popular idea in molecular biology is the geometric representation of a molecule by filling the space around the atom centers. It is assumed that the locations of the centers in space are known. There are three space filling models introduced in seminal papers by Lee and Richards [13], [19]. The *van der Waals* or *VW model* represents each atom, A , by a spherical ball of points at distance at most the van der Waals radius, r_A , from the center. Different types of atoms may have different radii. Overlapping spheres are indicative of chemically binding. The interaction of the VW model with a solvent sphere of radius $\alpha > 0$ is captured by the *solvent accessible* or *SA model*. It grows the sphere representing A to radius $r_A + \alpha$. With this amount of growth we have a solvent sphere overlapping the VW model iff its center belongs to the SA model.

The third is the *molecular surface* or *MS model*. In recognition of the first widely available computer program constructing it [4] the boundary of the MS model is sometimes referred to as the *Connolly surface*. Intuitively, the MS model is obtained by rolling the solvent sphere about the VW model and filling in the space inaccessible to that sphere. During the rolling motion, the solvent sphere stays in contact with the VW surface and its center sweeps out the SA surface. The molecular surface consists of patches that correspond to faces of the SA model. We have sphere patches, torus patches, and inverse sphere patches that correspond to sphere patches, circle arcs, and corner points of the SA model.

The two major advantages of the MS over the SA model are apparent smoothness and faithful size representation. Unfortunately, the molecular surface can have self-intersections resulting in a variety of ambiguities. We can clip at places of self-intersection and thus create a nonsmooth surface bounding an unambiguous subset of space. The topology of this subset is usually not the same as that of the SA model. Furthermore, it is difficult to construct the clipped surface automatically in a reliable manner and no robust software exists today. Alternatively, we can leave the surface unclipped. This helps in the robust construction from the face complex of the SA model [1]. Observe that the unclipped surface is still not smooth: self-intersecting torus patches connect to inverse sphere patches in sharp forward and backward folds.

Molecular Skin Surface. We advocate the body defined in Section 4 as a new space filling model that is geometrically unambiguous and otherwise similar to the MS model:

1. The body expresses the interaction with a solvent sphere without inflating atoms.

2. The skin or surface of the body does not intersect itself.
3. The skin is tangent continuous.

How exactly do we define and construct the body representing a molecule? First, we settle on a shrinking factor $s = \frac{1}{2}$. This choice is suggested by considerations of complementarity described below. An atom A with center z_A and van der Waals radius r_A is represented by the sphere $p_A = (z_A, \sqrt{2}r_A)$. Let P denote the resulting set of spheres. The *skin model* of the molecule and its surface are

$$\begin{aligned}\text{bdy } P &= \text{bdy}^{1/2} P, \\ \text{skn } P &= \text{bd bdy } P.\end{aligned}$$

Observe that the union of balls bounded by spheres in P is in general not the same as the SA model. An extension of the skin concept where different spheres are reduced by different amounts might be a worthwhile research project. The shrinking factor along edges, triangles, and tetrahedra of spheres can be varied by linear interpolation.

Docking. The docking problem for molecules is a question of shape and energy complementarity. We only comment on shape complementarity. Most important in this context is Theorem 14: for each molecule $M = \text{bdy } P$ there exists another skin model $L = \text{bdy } Q$ complementary to M and with perfectly matching skin. L is unique and given by the orthogonal complement Q of P .

There are several reasons why L falls short of modeling an actual molecule. The first is that L is unbounded because $L \cup M = \mathbb{R}^3$ and M is bounded. The second is that spheres in Q can have arbitrary size. Still, the existence and uniqueness of the set of spheres $Q = \text{orth } P$ with a complementary skin model is intriguing. It gives a mathematical justification for the heuristic construction of roughly complementary sphere packings in the approach to docking by Kuntz [12] and coworkers. More generally, it opens new ways of thinking about molecular complementarity. For example, notions of partial and of imperfect shape complementarity can be related to aspects of the difference between L and the skin model of a ligand.

Motion and Continuity. An actual molecule is constantly in motion and the distances and angles between its atoms vary with time. The study and simulation of this motion is the subject of molecular dynamics, which is a subfield of molecular modeling. Large numerical software packages exist that aim at the reliable simulation of one or several molecules over time, see, e.g., [15]. These packages produce sequences of discrete snapshots without continuous evolution between them. Depending on the amount and size of the change the reconstruction of such an evolution ranges from seemingly obvious to ambiguous and impossible.

A related example where motion and continuity comes up is the comparison of different observations of the same molecule. Let M be a protein and let M' be the same protein, maybe with a side chain removed. Models of M and M' are typically developed from independent crystalizations. The removal of the side chain triggers a rearrangement and repacking of atoms in its neighborhood. Resulting differences between M and M' are naturally described in causal or temporal terminology: “the side chain removal enlarges one of the voids,” etc. This description presupposes a correspondence between the voids

of M and M' . To claim a void is enlarged we first need to establish the sameness of a void in M and one in M' . This can only be done with a framework that connects M and M' over time. Homotopies, isotopies, and the topology changing deformations of Morse theory come to mind. If we use skin models we have access to deformation methods as described in Section 8 and can create continuous connections necessary to argue about motion and its effects.

10. Discussion

This paper presents a new paradigm for surface design. It has its foundations in the combinatorial subdisciplines of geometry, topology, and algorithms. The theory is reasonably complete and answers the basic questions. There are many problems that remain and plenty of work that has to be done.

Generalizing Skin. The skin concept can be extended to allow different amounts of shrinking for different spheres. Consider the following framework. P is a finite set of spheres and $S: \text{conv } P \rightarrow [0, 1]$ assigns each sphere $p \in \text{conv } P$ its own shrinking factor, $S(p)$. The body and skin defined by P and S are

$$\begin{aligned} \text{bdy}^S P &= \bigcup_{p \in \text{conv } P} \text{ucl } p^{S(p)}, \\ \text{skn}^S P &= \text{bd } \text{bdy}^S P. \end{aligned}$$

Even in the most general case where the assignment is arbitrary and not even continuous, S preserves the homotopy type of the body. What additional properties do we gain if we require S to be continuous? Call S *complementable* if there is an assignment $T: \text{conv } Q \rightarrow [0, 1]$, $Q = \text{orth } P$, with $\text{skn}^S P = \text{skn}^T Q$. Is there a characterization of complementable assignments in terms of local bounds on curvature?

Rendering and Triangulating Skin. This paper claims that skin has the potential to be used in a wide range of modeling applications requiring surface design and deformation. To live up to that claim it is important that fast and robust software be produced. What exactly this software constructs depends on the targeted application. Among the candidates, we single out the direct representation of patches and approximation through triangulations.

If the sole purpose of the construction is the graphical display of skin we can take advantage of available hardware that renders surfaces directly. The hardware supports low-degree algebraic surfaces, which includes degree-2 skin patches. There is a difficulty caused by complicated patch boundaries, which needs to be handled by surface trimming. Each patch is the intersection of a sphere or a hyperboloid with a convex polyhedron. It is possible that the intersection consists of several components, and components may have holes. Furthermore, there is no a priori upper bound on the number of boundary edges of even a single patch component.

Scientific applications require a representation supporting numerical computations. A typical such representation is a piecewise linear surface approximating the skin. We

have a *triangulation* if all linear pieces are triangles and the surface is homeomorphic to the skin. The construction of a triangulation is helped by the tangent continuity of skin. We propose developing an algorithm that adjusts the vertex density to the local curvature. Local point rearrangements can be used to get angles close to 60° . Theorem 18 is used to guarantee the approximating surface is homeomorphic to the skin by guiding decisions about constructing and fitting triangles. Without combinatorial guidance these decisions create robustness problems at or close to degenerate configurations.

Connecting Skin over Time. Fact 17 can be extended to more general pairs of skin. Suppose P_0 and P_1 are finite sets of spheres and we construct a map $f: \mathbb{R}^3 \rightarrow \mathbb{R}$ with

$$\begin{aligned} \text{skn}^{1/2} P_0 &= f^{-1}(0), \\ \text{skn}^{1/2} P_1 &= f^{-1}(1), \end{aligned}$$

see Section 8. The continuous sequence of preimages $f^{-1}(\tau)$, $\tau \in [0, 1]$, defines a deformation from one skin to the other. To what extent is it possible to establish a relationship between the points of the $f^{-1}(\tau)$?

Suppose first that the sequence of preimages avoids all critical points of f . Then there is a one-parameter family of diffeomorphisms between the preimages [14]. An interval of such diffeomorphisms forms an isotopy. What assumptions on the map f are necessary to construct the isotopy algorithmically? If the interval of preimages contains critical points of f , then there is no hope of constructing an isotopy because we get preimages of different homotopy types. Is there a weaker notion of a map that is almost an isotopy in the sense that it violates the conditions only locally around the critical points?

References

1. N. Akkiraju, H. Edelsbrunner, P. Fu, and J. Qiang. Viewing geometric protein structures from inside a CAVE. *IEEE Comput. Graphics Appl.* **16** (1996), 58–61.
2. H.-L. Cheng, H. Edelsbrunner, and P. Fu. Shape space from deformation. Report rgi-tech-98-011, Raindrop Geomagic, Champaign, Illinois, 1998.
3. S.-W. Cheng, H. Edelsbrunner, P. Fu, and K.-P. Lam. Design and analysis of planar shape deformation. In *Proc. 14th Ann. Sympos. Comput. Geom.*, 1998, to appear.
4. M. L. Connolly. Analytical molecular surface calculation. *J. Appl. Cryst.* **16** (1983), 548–558.
5. B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR Otdel. Mat. Est. Nauk* **7** (1934), 793–800.
6. C. J. A. Delfinado and H. Edelsbrunner. An incremental algorithm for Betti numbers of simplicial complexes on the 3-sphere. *Comput. Aided Geom. Design* **12** (1995), 771–784.
7. H. Edelsbrunner. The union of balls and its dual shape. *Discrete Comput. Geom.* **13** (1995), 415–440.
8. H. Edelsbrunner. Smooth surfaces for multi-scale shape representation. In *Proc. 15th Conf. Found. Software Techn. Theoret. Comput. Sci.* 1995, Lecture Notes Comput. Sci., Vol. 1026, Springer-Verlag, Berlin, pp. 391–412.
9. H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics* **13** (1994), 43–72.
10. G. Farin. *Curves and Surfaces for Computer Aided Geometric Design*. Academic Press, Boston, Massachusetts, 1988.
11. V. Guillemin and A. Pollack. *Differential Topology*. Prentice-Hall, Englewood Cliffs, New Jersey, 1974.
12. I. W. Kuntz. Structure-based strategies for drug design and discovery. *Science* **257** (1992), 1078–1082.
13. B. Lee and F. M. Richards. The interpretation of protein structures: estimation of static accessibility. *J. Mol. Biol.* **55** (1971), 379–400.

14. J. Milnor. *Morse Theory*. Annals Math. Studies, Princeton University Press, Princeton, New Jersey, 1963.
15. M. Nelson, W. Humphrey, A. Gürsoy, A. Dalke, L. V. Kalé, R. Skeel, K. Schulten, and R. Rufkin. MDScope—a visual computing environment for structural biology. *Comput. Phys. Commun.* **91** (1995), 111–134.
16. D. Pedoe. *Geometry: a Comprehensive Course*. Dover, New York, 1988.
17. R. E. Pfeifer and C. van Hook. Circles, vectors, and linear algebra. *Math. Mag.* **66** (1993), 75–86.
18. H. Qin and D. Terzopoulos. D-nurbs: a physics-based geometric design framework. *IEEE Trans. Visual. Comput. Graphics* **2** (1996), 85–96.
19. F. M. Richards. Areas, volumes, packing, and protein structures. *Ann. Rev. Biophys. Bioengr.* **6** (1977), 151–176.
20. G. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. *J. Reine Angew. Math.* **133** (1907), 97–178.

Received December 12, 1996, and in revised form December 4, 1997.