Weak Witnesses for Delaunay triangulations of Submanifold

Dominique Attali, Herbert Edelsbrunner, Yuriy Mileyko

Weak Witnesses for Delaunay Triangulations of Submanifolds *

Dominique Attali† LIS, INPG-CNRS
Herbert Edelsbrunner‡ Duke University and Geomagic
Yuriy Mileyko§ Duke University

Abstract

The main result of this paper is an extension of de Silva's Weak Delaunay Theorem to smoothly embedded curves and surfaces in Euclidean space. Assuming a sufficiently fine sampling, we prove that $i + 1$ points in the sample span an i-simplex in the restricted Delaunay triangulation iff every subset of the $i + 1$ points has a weak witness.

CR Categories: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—[Geometrical problems and computations, Computations on discrete structures]; I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—[Curve, surface, solid, and object representations]

Keywords: Computational geometry and topology, curves, surfaces, restricted Delaunay triangulations, witness complexes

1 Introduction

This paper contributes to the growing literature on extracting information from sampled point data. In particular, we are interested in shape reconstruction for data distributed in low-dimensional subspaces of ambient space.

Motivation. The broad availability of powerful hardware drives the emergence of data analysis as new paradigm in many areas of science and engineering. We now routinely collect large amounts of data, challenging our ability to extract relevant information fast enough and in a meaningful format. It is common to interpret the data items as points in some Euclidean space. The data as a whole is referred to as a *point cloud*, which emphasizes that we deal with large numbers and require analysis methods that summarize and simplify without losing sight of important details that may be hidden within the wealth of measurements.

To bring order into the various types of data analysis questions, we lay them out on an axis from coarse to fine. An example of a coarse analysis is the decomposition into clusters. In this paper we are interested in the fine end of the spectrum and in particular in the reconstruction of shapes from point clouds. An example is the reconstruction of a geometric shape from 3D scan data consisting of points measured on the surface of a physical object. We have the points distributed on a two-dimensional subspace of three-dimensional Euclidean space. That the dimension of the ambient space eclipses the intrinsic dimension of the data is typical [Amenta et al. 2007; Attali et al. 2003] and poses challenges as well as opportunities in the analysis. We need dimension reduction techniques but also methods that adapt to the intrinsic rather than the ambient dimension of the data.

Prior work and results. The problem of reconstructing shapes from point clouds in three-dimensional Euclidean space has been studied in computer graphics [Bernardini et al. 1999; Carr et al. 2001; Hoppe et al. 1992], computational geometry [Amenta and Bern 1999; Chazal et al. 2006; Chazal and Lieutier 2006], and other areas [Niyogi et al. to appear; Zhao et al. 2001]. The typical approach in computational geometry starts with the (three-dimensional) Delaunay triangulation and aims at extracting the (two-dimensional) restricted Delaunay triangulation. Assuming we know the surface from which the data points are sampled, the restricted Delaunay triangulation consists of all simplices whose dual Voronoi cells have non-empty intersections with the surface [Edelsbrunner and Shah 1997]. Since the surface is generally unknown, the algorithms often substitute constructs derived from the data or iterate until the reconstructed surface is its own restricted Delaunay triangulation [Edelsbrunner 2004; Giesen and John 2002].

The desire to free oneself from the ambient dimension motivates the introduction of witness complexes in [de Silva and Carlsson 2004]. Building on the work in [Martinetz and Schulten 1994], de Silva and Carlsson distinguish between two kinds of data points, *landmarks* used in the construction of a complex and *witnesses* used to guide the selection of simplices connecting the landmarks. The method is based on the intuition that a large cloud of witnesses gives a good representation of the subspace occupied by the data and that relatively few landmarks are needed to give a satisfactory reconstruction. A crucial concept in this approach is the notion of a *weak witness* of $i + 1$ landmarks, which is a point for which these are the $i + 1$ nearest landmarks. The link to the earlier work is provided by de Silva who proves that the $i + 1$ landmarks span an i-simplex in the Delaunay triangulation iff each of its subsets has a weak witness [de Silva 2003]. This implies that the witness complex approximates the Delaunay triangulation and reaches it in the limit, when every point of the ambient space becomes a witness.

In this paper, we extend de Silva's result to submanifolds of Euclidean space. Assuming a sufficiently fine sampling of landmarks on a smoothly embedded curve or surface, we prove that the witness complex approximates the restricted Delaunay triangulation which it reaches in the limit, when every point of the submanifold becomes a witness. This result does not depend on the ambient dimension. We contrast this structural theorem to a recent result by Guibas and Oudot which requires that the landmark are placed according to a particular strategy and that the dimensions of the submanifold and the ambient space differ by exactly by one [Guibas and Oudot 2007]. Extensions to our work are discussed in Section 6.

Outline. Section 2 gives a detailed statement of our result. Section 3 presents basic geometric and topological tools. Sections 4 and 5 prove our result for curves and surfaces. Section 6 concludes the paper. Appendix A reviews basic differential geometry facts for smooth 2-manifolds.

*Research by the authors is partially supported by the Defense Advanced Research Projects Agency (DARPA) under grants HR0011-05-1-0007 and HR0011-05-1-0057 and by CNRS under grant PICS-3416.

†e-mail: Dominique.Attali@lis.inpg.fr
‡e-mail: edels@cs.duke.edu
§e-mail: yury@cs.duke.edu

2 Definitions and Result

In this section, we introduce the necessary definitions and give a complete description of our results. We begin with a review of the Weak Delaunay Theorem by de Silva to which our results are related.

In Euclidean space. Consider a finite set of points in d-dimensional Euclidean space, $L \subseteq \mathbb{R}^d$. We call these points *landmarks* because they are the vertices of a simplicial complex we are going to build. For the moment, we ignore the issue of geometric realizability and construct the complex abstractly, calling every non-empty subset $\sigma \subseteq L$ a *simplex*. Its *dimension* is one less than its cardinality, $\dim \sigma = \operatorname{card} \sigma - 1$. A *face* of σ is a non-empty subset $\tau \subseteq \sigma$. It is *improper* if $\tau = \sigma$ and otherwise *proper*. A *simplicial complex* is a collection of simplices that is closed under the face relation.

Whether or not we add a simplex to our complex depends on the position of its vertices among the other landmarks. To make this precise, we call a point $x \in \mathbb{R}^d$ a *weak (Delaunay) witness* of σ if $\|x-a\| \leq \|x-b\|$ for all $a \in \sigma$ and $b \in L - \sigma$. A *strong (Delaunay) witness* of σ is a weak witness that is equidistant from all vertices, $\|x-a\| = \|x-a'\|$ for all $a, a' \in \sigma$. A crucial difference between the two notions is that the weak witnesses of a simplex generally form a set with positive d-dimensional measure while the strong witnesses form a set of measure zero. It follows that the probability of finding a strong witness by sampling is zero. A fundamental result by de Silva says that the existence of weak witnesses can be used to infer the existence of strong witnesses.

WEAK DELAUNAY THEOREM [DE SILVA 2003]. Let L be a finite set in $\mathbb{R}^d$. If every face of a simplex $\sigma \subseteq L$ has a weak witness then σ has a strong witness.

We will give a proof of this result in Section 3. Similar to de Silva's our proof is constructive, giving a strong witness in the convex hull of the weak witnesses of the simplex and its faces.

On a submanifold. We are interested in the case in which the witnesses form a subset of Euclidean space and the landmarks are sampled from this subset. Specifically, we consider a *dimension k submanifold* $\mathbb{M}$ of $\mathbb{R}^d$ which, by definition, is a compact k-manifold without boundary that is smoothly embedded in d-dimensional Euclidean space. It is easy to see that de Silva's theorem does not hold for submanifolds. Specifically, we can have a simplex $\sigma \subseteq L \subseteq \mathbb{M}$ that has no strong witness on $\mathbb{M}$ even though all its faces have weak witnesses on $\mathbb{M}$. As suggested by the example in Figure 1, the implication fails because the sampling of the submanifold is not sufficiently fine.

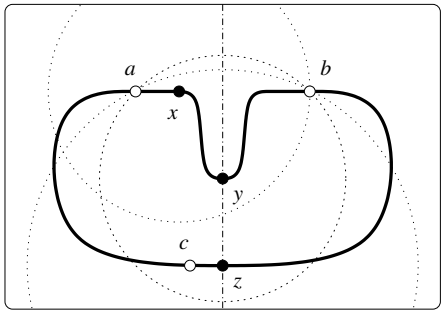

Figure 1: A dimension 1 submanifold of $\mathbb{R}^2$ with three landmarks, a, b, c. The point x is a weak witness of the edge $\{a, b\}$ but the only points y and z of the submanifold equidistant to a and b are closer to c and are therefore not strong witnesses of that edge.

Sampling condition. Since $\mathbb{M}$ is smooth and k-dimensional, the tangent space $\mathrm{T}\mathbb{M}_x$ at a point $x \in \mathbb{M}$ is a k-dimensional linear subspace of $\mathbb{R}^d$. For a non-zero tangent vector $v \in \mathrm{T}\mathbb{M}_x$, the *sectional curvature*, $\kappa(x, v)$, is the (absolute) curvature of a geodesic that passes through x in the direction v. We write $\kappa(x)$ for the local maximum, over all $v \in \mathrm{T}\mathbb{M}_x$, and

$$\kappa \quad = \quad \max_{x \in \mathbb{M}} \max_{v \in \mathrm{T}\mathbb{M}_x} \kappa(x, v)$$

for the globally *maximum (absolute) sectional curvature* of $\mathbb{M}$. Since $\mathbb{M}$ is smoothly embedded in $\mathbb{R}^d$, we can define the *normal space* $\mathrm{N}\mathbb{M}_x = \mathrm{T}\mathbb{M}_x^{\perp}$ at a point $x \in \mathbb{M}$ consisting of all vectors $u \in \mathbb{R}^d$ that are orthogonal to all tangent vectors $v \in \mathrm{T}\mathbb{M}_x$. Since $\mathrm{T}\mathbb{M}_x$ has dimension k, $\mathrm{N}\mathbb{M}_x$ has dimension $d-k$. For each non-zero $u \in \mathrm{N}\mathbb{M}_x$, we let $\varrho(x, u)$ be the supremum of the radii r such that the open d-ball with center $x + r\frac{u}{\|u\|}$ and radius r has an empty intersection with $\mathbb{M}$. We write $\varrho(x)$ for the local minimum, over all $u \in \mathrm{N}\mathbb{M}_x$, referred to as the *local reach* of $\mathbb{M}$ at x, and

$$\varrho \quad = \quad \min_{x \in \mathbb{M}} \min_{u \in \mathrm{N}\mathbb{M}_x} \varrho(x, u)$$

for the global minimum, referred to as the *(global) reach* of $\mathbb{M}$. We note that $\varrho(x) \leq 1/\kappa(x)$ at every point $x \in \mathbb{M}$ and therefore $\kappa\varrho \leq 1$; see also [Niyogi et al. to appear].

DEFINITION. Given $\varepsilon > 0$, we call a finite set $L \subseteq \mathbb{M}$ an *ε-sample* of $\mathbb{M}$ if for every point $x \in \mathbb{M}$ there are at least $k+1$ landmarks $a \in L$ whose Euclidean distance to x is $\|x-a\| < \varepsilon\varrho$.

Note that ε-samples can be finer but not coarser than required. In other words, every ε'-sample with $\varepsilon' \leq \varepsilon$ is also an ε-sample of $\mathbb{M}$.

Main result. We are now ready to state our main result.

WEAK RESTRICTED DELAUNAY THEOREM. For $k = 1, 2$ there exists $\varepsilon > 0$ such that for every ε-sample L of a dimension k submanifold $\mathbb{M}$ of $\mathbb{R}^d$, the following implication holds: if every face of a simplex $\sigma \subseteq L$ has a weak witness on $\mathbb{M}$ then σ has a strong witness on $\mathbb{M}$.

We say that a subset L of a submanifold $\mathbb{M}$ of $\mathbb{R}^d$ has the *weak Delaunay property* if every simplex $\sigma \subseteq L$ whose faces all have weak witnesses on $\mathbb{M}$ has a strong witness on $\mathbb{M}$. Let $\varepsilon_{k,d}$ be the supremum of all values of ε such that every ε-sample of every dimension k submanifold of $\mathbb{R}^d$ has the weak Delaunay property. We call $\varepsilon_k = \inf_{d>k} \varepsilon_{k,d}$ the *weak Delaunay constant* of dimension k submanifolds. Our result says that for $k = 1, 2$ the weak Delaunay constant of dimension k submanifolds is positive, that is, $\varepsilon_k > 0$. For curves we establish tight upper and lower bounds giving $\varepsilon_1 = \sqrt{3}$. For surfaces we show $\frac{1}{\sqrt{5}} \leq \varepsilon_2 \leq \sqrt{2}$, leaving a substantial gap between the two bounds.

3 Preliminaries

In this section, we prepare the proof of our main result. Specifically, we give a new proof of de Silva's original theorem and we introduce a basic topological lemma that will allow us to adapt this proof for submanifolds.

Euclidean space. Here we give a proof of de Silva's Weak Delaunay Theorem stated in Section 2. We proceed by induction over the dimension of the simplices. The vertices (0-simplices) obviously satisfy the claim and form the induction basis.

Let σ be a simplex of dimension i whose faces all have weak witnesses. By induction hypothesis, every proper face of σ has a strong witness. Let x_0 be a weak witness of σ and let $\tau \subseteq \sigma$ be the face spanned by the subset of landmarks in σ that are furthest and at equal distance from x_0. If $\tau = \sigma$ then x_0 is equidistant from all $i+1$ vertices and therefore a strong witness of σ. Else τ is a proper face with a strong witness x_1. We interpolate between the two witnesses by defining $x_s = (1-s)x_0 + sx_1$. The vertices of τ lie on the $(d-2)$-sphere common to the two $(d-1)$-spheres centered at x_0 and x_1 that both contain τ. For each $s \in [0,1]$ consider the closed ball B_s with center x_s whose bounding sphere passes through the landmarks in τ. The landmarks in $\sigma - \tau$ lie in the interior of B_0 and outside the interior of B_1. For intermediate values of s, they can lie in the interior, on the boundary, or outside B_s. By construction, the sphere ∂B_s passes through $\partial B_0 \cap \partial B_1$. Since x_s lies between x_0 and x_1, this implies $B_s \subseteq B_0 \cup B_1$ for all $s \in [0,1]$, as illustrated in Figure 2. The ball can therefore not pick up any

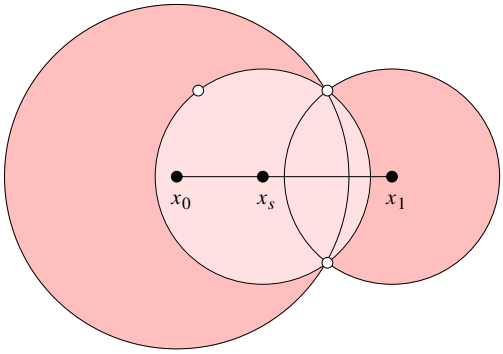

Figure 2: The circle around x_s passes through the intersection points of the circles around x_0 and x_1. It also passes through the landmark inside the circle around x_0.

new landmarks as its center moves from x_0 to x_1, only lose some. We stop the motion at the smallest value t of s for which one of the landmarks in $\sigma - \tau$ escapes from the interior to the boundary of B_s. This value t exists because eventually, for $s = 1$, all landmarks will have escaped from the interior. The new point x_t is either a strong witness or another weak witness of σ. In the latter case, the number of furthest among the $i+1$ vertices of σ increased by at least one. We can therefore repeat the linear interpolation, substituting x_t for x_0. After at most i steps, all $i+1$ vertices of σ are furthest from the stopping point of the interpolation, implying we have arrived at the strong witness, whose existence has thus been established. This completes the proof of de Silva's Weak Delaunay Theorem.

REMARK. The inductive step in the above proof works equally well for simplices of dimension $i \leq d$ and $i > d$. Since each inductive step increases the dimension of the lowest-dimensional sphere that contains τ, we need at most d steps to arrive at the strong witness, even if $i > d$. Also note that the proof constructs a strong witness which is a convex combination of the weak witnesses of σ and of its faces.

A topological lemma. We now state and prove a refinement of Lemma 7 in [Amenta and Bern 1999] that will allow us to turn the Euclidean spheres in $\mathbb{R}^d$ into topological spheres in the submanifold. With this key ingredient, we will be able to mimic the Euclidean argument on the submanifold.

LOCAL REACH LEMMA. Let $\mathbb{M}$ be a dimension k submanifold of $\mathbb{R}^d$, B a closed ball, and ϱ_B the minimum of $\varrho(x)$ over all points $x \in B \cap \mathbb{M}$. If the center of B is at distance $\delta < \varrho_B$ from $\mathbb{M}$ and the radius r satisfies $\delta < r < 2\varrho_B - \delta$ then $B \cap \mathbb{M}$ is a topological k-ball.

PROOF. Let z be the center of B and let $f : \mathbb{M} \to \mathbb{R}$ be defined by $f(x) = \|x - z\|^2$. The intersection of $\mathbb{M}$ with the ball is the sublevel set defined by the radius, $B \cap \mathbb{M} = f^{-1}[0, r^2]$. If this sublevel set contains only one critical point then this is a minimum, with function value $\delta^2 < r^2$, and the sublevel set is a topological ball whose dimension is the same as that of the submanifold [Milnor 1963]. Else there are at least two critical points, including $x \neq y$ with $\delta^2 = f(x) \leq f(y) \leq r^2$. Since y is critical, $u = z - y$ is a non-zero normal vector of $\mathbb{M}$ at y. Consider the closed d-ball B' with radius $r' = \frac{1}{2}(r + \delta)$ and center $y' = y + r'\frac{u}{\|u\|}$. As

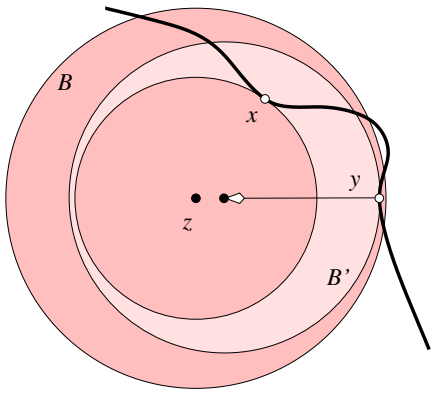

Figure 3: The three balls are nested and all contain the point x closest to the center of B.

illustrated in Figure 3, B' contains the ball with center z and radius δ and therefore also the point $x \in \mathbb{M}$. But since $r' < \varrho_B \leq \varrho(y)$, x belongs to the interior of the d-ball with radius $\varrho(y)$ and center $y + \varrho(y)\frac{u}{\|u\|}$. This contradicts that $\varrho(y)$ is the local reach of $\mathbb{M}$ at y. ∎

REMARK. We get the strictly weaker (global) Reach Lemma by substituting ϱ for ϱ_B in the Local Reach Lemma. In many but not all cases, this weaker statement will suffice for our purposes.

4 Curves

In this section, we prove the Weak Restricted Delaunay Theorem for a closed curve $\mathbb{M}$ smoothly embedded in $\mathbb{R}^d$. The proof is relatively straightforward and we are able to give matching upper and lower bounds for the required sampling density. Specifically, we prove that for $\varepsilon = \sqrt{3} = 1.732\ldots$ every edge of an ε-sample that has a weak witness on $\mathbb{M}$ also has a strong witness on $\mathbb{M}$ and that no triangle has three edges each of which has a weak witness on $\mathbb{M}$.

Edges. The main technical ingredient is the 1-dimensional version of the Local Reach Lemma. To use it, we write B_x for the smallest closed d-ball with center $x \in \mathbb{M}$ that contains at least two landmarks. By definition of ε-sample, the radius of B_x is $r < \sqrt{3}\varrho$. Since this is less than 2ϱ, the Local Reach Lemma implies that B_x intersects $\mathbb{M}$ in an interval (a closed topological 1-ball). If x is a weak but not a strong witness of the edge $\{a, b\}$ then B_x contains one landmark in the interior and the other on the boundary, as in Figure 4. Since B_x intersects $\mathbb{M}$ in a single interval it contains the entire arc from a to b, implying that this arc does not contain any other landmarks. We let y be a point on this arc that is equidis-

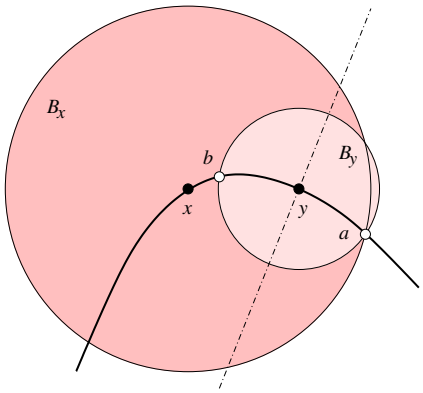

Figure 4: The point x is a weak witness of $\{a, b\}$ and the point y is a strong witness of this edge.

tant from a and b. Because $B_y \cap \mathbb{M}$ is an interval that contains at least two landmarks, the only possibility is that this interval is the arc from a to b. In particular, no other landmarks lies inside B_y, showing that y is a strong witness of the edge $\{a, b\}$.

Triangles. To prove that no triangle has three edges each with a weak witness, we show that the sampling condition implies there are at least four landmarks on each component of $\mathbb{M}$. In this case, each triangle has at least one edge whose landmarks are not contiguous along $\mathbb{M}$. Since the ball defined by a weak witness of this edge meets $\mathbb{M}$ in an interval, it contains at least one additional landmark in its interior, contradicting the definition of weak witness.

To show that there are at least four landmarks per component, we assume that $\mathbb{M}$ is connected and contains only three landmarks, a, b, c, decomposing $\mathbb{M}$ into three arcs, ab, bc, ca. Assuming ab is the shortest, we concatenate the two other arcs to get $C = bc \cup ca$. Since $1/\varrho$ is an upper bound on the curvature of $\mathbb{M}$ at every point, the length of $\mathbb{M}$ is at least $2\pi\varrho$, the length of the circle with radius ϱ. It follows that C has length at least $\frac{4}{3}\pi\varrho$. The midpoint, x, decomposes C into two arcs of length at least $\frac{2}{3}\pi\varrho$ each. Consider the closed d-ball with center x and radius $\sqrt{3}\varrho$. It contains precisely two thirds of any circle with radius ϱ that passes through its center. The Reach Lemma implies that the d-ball intersects $\mathbb{M}$ in a single arc $A \subseteq \mathbb{M}$. The point x decomposes A into two pieces, and because $1/\varrho$ is an upper bound on the curvature, each piece has length at most $\frac{2}{3}\pi\varrho$, the length of the two arcs in which the circle of radius ϱ intersects the d-ball. Hence $A \subseteq C$, which implies that A contains at most one landmark in its interior, namely c. This contradicts that a, b, c form an ε-sample, for $\varepsilon = \sqrt{3}$, and thus implies that $\mathbb{M}$ contains at least four landmarks, as required.

Upper bound. We note that $\varepsilon = \sqrt{3}$ is tight. Indeed, for every $\varepsilon > \sqrt{3}$ we can construct a counterexample consisting of a circle with three landmarks placed at the vertices of an inscribed equilateral triangle. The landmarks form an ε-sample of the circle, the triangle has a weak witness, all three edges have strong witnesses, but the triangle itself does not have a strong witness on the circle. This implies that the weak Delaunay constant for dimension 1 submanifolds is $\varepsilon_1 = \sqrt{3}$.

5 Surfaces

In this section, we prove the Weak Restricted Delaunay Theorem for surfaces. After setting the stage, we consider edges, triangles, and higher-dimensional simplices, in this order.

Topological spheres. Let $\mathbb{M}$ be a dimension 2 submanifold of $\mathbb{R}^d$, that is, a compact closed surface smoothly embedded in d-dimensional Euclidean space. According to the Reach Lemma, a closed d-ball whose center lies on $\mathbb{M}$ and whose radius is less than 2ϱ intersects $\mathbb{M}$ in a topological disk. It follows that the bounding $(d-1)$-sphere intersects $\mathbb{M}$ in a topological circle. In preparation of the proof of our main result for surfaces, we now consider two d-balls but limit their radii to less than $\varepsilon\varrho$, for $\varepsilon = 1/\sqrt{5}$. Let B_x be a d-ball with center $x \in \mathbb{M}$ and radius $r_x < \varepsilon\varrho$, and similarly let B_y be a d-ball with center $y \in \mathbb{M}$ and radius $r_y < \varepsilon\varrho$. Excluding $B_x = B_y$ as a possibility, the two bounding $(d-1)$-spheres are disjoint, meet at a point, or intersect in a $(d-2)$-sphere. Assuming the last case, let B_{xy} be the $(d-1)$-ball whose boundary is that $(d-2)$-sphere. The center of B_{xy} is not necessarily on $\mathbb{M}$ and its radius is $r_{xy} \leq \min\{r_x, r_y\} < \varepsilon\varrho$.

INTERVAL LEMMA. The intersection of B_{xy} with $\mathbb{M}$ is either empty, a point, or a closed topological interval.

PROOF. Let P_{xy} be the $(d-1)$-plane that contains B_{xy}. By construction, $v = y - x$ is normal to P_{xy}. Assuming $B_{xy} \cap \mathbb{M}$ is nonempty, let z be a point in this intersection and note that $\|x - z\| < \varepsilon\varrho$. Recall that κ is the maximum sectional curvature, over all points of $\mathbb{M}$ and all tangent directions, and that $\kappa\varrho \leq 1$. By Proposition 6.3 in [Niyogi et al. to appear], the geodesic distance between the two points x and z is $d(x, z) \leq \varrho(1-\sqrt{1-2\varepsilon})$. For $\varepsilon = 1/\sqrt{5}$ this gives $d(x, z) < \frac{\pi}{\kappa}$ which justifies the application of Property I in Appendix A, improving the bound to $d(x, z) < \frac{2}{\kappa}\arcsin\frac{\varepsilon}{2}$. Similarly, $\|x - y\| < 2\varepsilon\varrho$ and $d(x, y) \leq 2d(x, z) < \frac{\pi}{\kappa}$ and therefore $d(x, y) < \frac{2}{\kappa}\arcsin\varepsilon$. Recall that the angle between v and the tangent plane at z is the minimum angle between v and a vector $u \in \mathrm{T}\mathbb{M}_z$. To bound this angle, we use the triangle inequality followed by Properties III, II, I, in this sequence,

$$\begin{aligned} \angle v\mathrm{T}\mathbb{M}_z &\leq \angle v\mathrm{T}\mathbb{M}_x + \angle \mathrm{T}\mathbb{M}_x\mathrm{T}\mathbb{M}_z \\ &\leq \frac{\kappa}{2}d(x, y) + \kappa d(x, z) \\ &< \arcsin\varepsilon + 2\arcsin\frac{\varepsilon}{2}. \end{aligned}$$

Because the arcsin function is convex, the last line is bounded from above by $\arcsin 2\varepsilon$ which for $\varepsilon = 1/\sqrt{5}$ is less than $\frac{\pi}{2}$. Generically, $P_{xy} \cap \mathbb{M}$ is a collection of curves, and since the angle between v and the tangent plane at a point z is bounded away from $\frac{\pi}{2}$, the intersection is a curve in a neighborhood of z, even in the non-generic case. To get a lower bound on the local reach of this curve at $z \in B_{xy}$, we let the unit vector $u \in \mathrm{T}\mathbb{M}_z$ minimize the angle with v and write $\varphi = \angle vu = \angle v\mathrm{T}\mathbb{M}_z$. Decomposing v into its components in $\mathrm{T}\mathbb{M}_z$ and $\mathrm{N}\mathbb{M}_z$, we get $v = u\cos\varphi + w\sin\varphi$. The two $(d-1)$-spheres with centers $z \pm \varrho w$ and radius ϱ are both tangent to $\mathbb{M}$ at z and sandwich the manifold between them. Cutting them with the $(d-1)$-plane P_{xy} we get two $(d-2)$-spheres of radius $\varrho\cos\varphi$ each. They are both tangent to $P_{xy} \cap \mathbb{M}$ at z and sandwich the curve between them. This implies that the local reach of the curve at z is

$$\begin{aligned} \rho(z) &\geq \varrho\cos\angle v\mathrm{T}\mathbb{M}_z \\ &> \varrho\cos(\arcsin 2\varepsilon) \\ &= \varrho\sqrt{1-4\varepsilon^2}. \end{aligned}$$

Since $\varepsilon = 1/\sqrt{5}$ we get $\rho(z) \geq \varepsilon\varrho$. But this exceeds the radius of the $(d-1)$-ball, $r_{xy} < \rho(z)$. Applying the Local Reach Lemma to the cross-section within P_{xy} implies that B_{xy} meets $\mathbb{M}$ either at a point, namely if its center is at distance r_{xy} from $\mathbb{M}$, or in a closed topological interval. □

Edges. We now consider the first case in the inductive proof of the Weak Restricted Delaunay Theorem. Specifically, we show that for $\varepsilon = 1/\sqrt{5}$ every edge that has a weak witness on $\mathbb{M}$ also has a strong witness on $\mathbb{M}$.

Let $x_0 \in \mathbb{M}$ be a weak witness of the edge $\{a, b\}$, assume $\|x_0 - b\| < \|x_0 - a\|$, and let B_0 be the d-ball with center x_0 and radius $\|x_0 - a\| < \varepsilon\varrho$, as shown in Figure 5. To construct a strong witness, we move the center along a particular path $\alpha : [0, 1] \to \mathbb{M}$ that connects $x_0 = \alpha(0)$ with $a = \alpha(1)$. To construct the path, we let D be the *diameter d-ball* of x_0 and a, with center $\frac{1}{2}(x_0 + a)$ and radius $\frac{1}{2}\|x_0 - a\|$, and we let P be the $(d-1)$-plane passing through x_0 and a that contains all normal directions of $\mathbb{M}$ at a. The intersection, $D \cap P$, is a $(d-1)$-ball with radius less than $\frac{\varepsilon}{2}\varrho$. Using Properties I and II, we see that the angle between the tangent planes at a and at a point $z \in D \cap P \cap \mathbb{M}$ is $\angle \mathrm{T}\mathbb{M}_a \mathrm{T}\mathbb{M}_z \leq \kappa d(a, z) < 2 \arcsin \frac{\varepsilon}{2}$. Since P contains $a + \mathrm{N}\mathbb{M}_a$, its normal vector, v, is contained in the tangent plane at a, $\angle v \mathrm{T}\mathbb{M}_a = 0$. By the triangle inequality, the angle between v and the tangent plane at z thus satisfies $\angle v \mathrm{T}\mathbb{M}_z < 2 \arcsin \frac{\varepsilon}{2} \leq \arcsin \varepsilon$. The local reach of the curve $P \cap \mathbb{M}$ at z is therefore $\rho(z) > \varrho \cos(\arcsin \varepsilon) = \varrho\sqrt{1 - \varepsilon^2}$. Since the radius of $D \cap P$ is less than that, the Local Reach Lemma implies that $D \cap P \cap \mathbb{M}$ is an interval. By construction, x_0 and a are the endpoints of this interval, and we let $\alpha : [0, 1] \to \mathbb{M}$ be a parametrization.

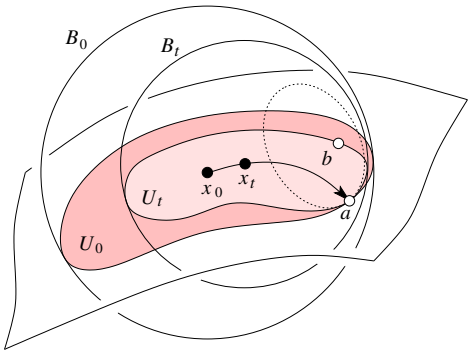

Figure 5: By choice of the path from x_0 to a, the topological disks are nested, with U_0 containing all others.

To describe the motion of the center, we write $x_s = \alpha(s)$, for $0 \leq s \leq 1$, and let B_s be the closed d-ball with center x_s and radius $\|x_s - a\| < \varepsilon\varrho$. Eventually, we run into $a = x_1$, which implies we must have passed a point x_t that is equidistant to a and b. To prove that x_t is a strong witness of the edge $\{a, b\}$ it suffices to show that $U_t = B_t \cap \mathbb{M}$ is contained in $U_0 = B_0 \cap \mathbb{M}$. By choice of ε, U_0 and U_t are both topological disks. The intersection of their boundaries is equal to the intersection of $\mathbb{M}$ with the $(d-2)$-sphere $\partial B_0 \cap \partial B_t$. By construction, the $(d-1)$-plane P passes through a and the centers of B_0 and B_s. Therefore, the normal vector v to P is tangent to the $(d-2)$-sphere $\partial B_0 \cap \partial B_t$ at point a. Since v is contained in $\mathrm{T}\mathbb{M}_a$ this $(d-2)$-sphere touches $\mathbb{M}$ tangentially at a and, by the Interval Lemma, it has a unique intersection point with $\mathbb{M}$, namely a. This implies that the boundaries of U_0 and U_t also intersect in a unique point, a, at which they touch each other tangentially, as in Figure 5. This leaves only two possibilities, namely $U_t \subseteq U_0$ or $U_0 \subseteq U_t$. To contradict the latter, we observe that $U_s \subseteq U_0$ for $s = 1$. Since U_s changes continuously while its boundary shares only the point a with ∂U_0, this property holds for all s and therefore also for $s = t$. It follows that x_t is a strong witness of the edge $\{a, b\}$, as desired.

Triangles. Next we show that for $\varepsilon = 1/\sqrt{5}$ every triangle whose faces all have weak witnesses on $\mathbb{M}$ also has a strong witness on $\mathbb{M}$.

Let $x_0 \in \mathbb{M}$ be a weak witness of $\{a, b, c\}$, assume $\|x_0 - c\|, \|x_0 - b\| < \|x_0 - a\|$, and let B_0 be the closed d-ball with center x_0 and radius $\|x_0 - a\| < \varepsilon\varrho$. The first step is the same as for an edge, namely moving the center toward a until its distance to b or to c is the same as that to a. The second step is similar to the first but different in detail. To emphasize the similarities, we reuse notation writing x_0 for the starting point of the second step which is the point x_t constructed in the first step. Without loss of generality, we assume $\|x_0 - c\| < \|x_0 - b\| = \|x_0 - a\|$, as in Figure 6. To construct a strong witness for the triangle, we move the center along a path $\alpha : [0, 1] \to \mathbb{M}$ connecting $x_0 = \alpha(0)$ with the strong witness $x_1 = \alpha(1)$ of the edge $\{a, b\}$, which exists inductively. To construct this path, let r be the larger of the two distances $\|x_0 - a\| = \|x_0 - b\|$ and $\|x_1 - a\| = \|x_1 - b\|$. Let B_{ab} be the $(d-1)$-ball of points whose distance from a and b is the same and at most r. Note that x_0 and x_1 both belong to B_{ab}. By the Interval Lemma, $B_{ab} \cap \mathbb{M}$ is a closed, topological interval. We let $\alpha : [0, 1] \to \mathbb{M}$ be a parametrization of the subinterval from x_0 to x_1 and write $x_s = \alpha(s)$.

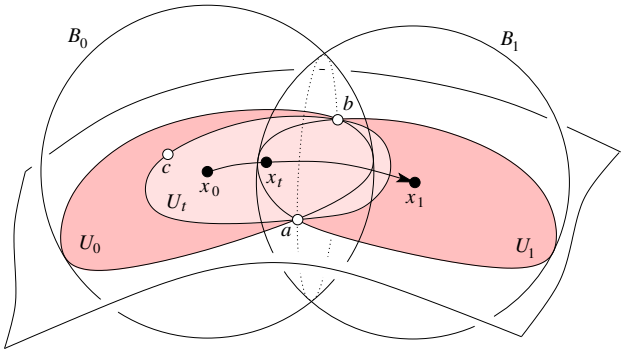

Figure 6: By choice of the path from x_0 to x_1 we get a pencil of topological disks, all contained in $U_0 \cup U_1$.

For the remainder of the analysis, let B_s be the smallest closed d-ball with center x_s that contains at least two landmarks. For example, B_0 contains c in its interior and a, b on its boundary. By definition of ε-sample, the radius of B_s is less than $\varepsilon\varrho$, for every s. The Reach Lemma thus implies that $U_s = B_s \cap \mathbb{M}$ is a topological disk, for every s. By the Interval Lemma, the boundaries of any two such topological disks intersect in a 0-sphere, that is, two points. For example, ∂U_0 and ∂U_1 intersect in points a and b. The two points decompose ∂U_i into two segments, and we call $\partial U_i \cap U_{1-i}$ the *inner segment* of ∂U_i, for $i = 0, 1$. For sufficiently small $s \geq 0$, ∂U_s passes through a and b and U_s contains the inner segment of ∂U_1. Symmetrically, for sufficiently large $s \leq 1$, ∂U_s passes through a and b and U_s contains the inner segment of ∂U_0. We use continuity to prove that

(i) ∂U_s passes through a and b, and

(ii) $U_s \subseteq U_0 \cup U_1$,

for every $s \in [0,1]$. By definition of B_s, U_s cannot contain a and b in its interior. Because ∂U_s intersects ∂U_0 and ∂U_1 in only two points each, U_s cannot pick up a new landmark point unless it first drops a and b from its boundary. But the latter is impossible unless U_s picks up a new landmark point first. This creates a deadlock situation and thus proves (i). We get (ii) because continuity now implies that all U_s contain the inner segments of ∂U_0 and ∂U_1. The only change during the movement of the center along α thus concerns the third landmark, c. It belongs to U_0 but not to U_1. Hence there is a value $t \in [0,1]$ such that ∂U_t passes through c. It follows that x_t is equidistant to a, b, c or, equivalently, that all three landmarks lie on ∂U_t, as in Figure 6. Since $U_t \subseteq U_0 \cup U_1$, it contains no other landmark, which implies that x_t is a strong witness of $\{a, b, c\}$, as desired.

Tetrahedra and beyond. Finally we show that for $\varepsilon = 1/\sqrt{5}$ every simplex of dimension three or higher whose faces all have weak witnesses on $\mathbb{M}$ has a strong witness on $\mathbb{M}$. Incidentally, the strong witnesses of all faces of dimension two or higher are the same.

Let $\{a, b, c, d\}$ be a tetrahedron whose faces all have weak witnesses. By induction hypothesis, its four triangles have strong witnesses, each defining a $(d-1)$-sphere intersecting $\mathbb{M}$ in a topological circle that passes through the three landmarks defining the triangle. If different, two such topological circles meet in at most two points, and because their triangles share two vertices, they intersect in these two shared landmarks. The landmarks decompose each topological circle into three arcs. We thus get a graph of four nodes (the landmarks) and twelve edges (the arcs). Since the circles meet only at landmarks, the edges do not cross, and because the circles are relatively small, the graph is drawn on a patch of $\mathbb{M}$ that is homeomorphic to a disk. In other words, we have a plane embedding of K_4, the complete graph with four nodes, in which every edge is doubled, as in Figure 7. But this implies that the middle node lies inside the topological circle passing through the outer three nodes, contradicting the construction of that circle around a strong witness of the triangle. The only resolution to this contradiction is that at least two of the topological circles are the same. This circle passes through all four nodes, forcing all four topological circles to be the same. In other words, the four triangles have a common strong witness, which is therefore also a strong witness of the tetrahedron. By the same argument, if an i-simplex with $i+1 > 4$ vertices has a weak witness for each face then all its faces have a common strong witness, which is also a strong witness of the i-simplex.

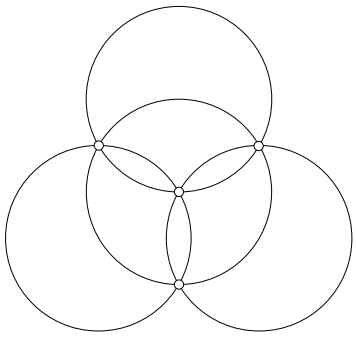

Figure 7: A crossing-free drawing of K_4, with edges doubled up.

Upper bound. Our lower bound for ε_2 is certainly not tight. Following a construction in [de Silva 2003], we get an upper bound by letting $\mathbb{M}$ be the unit 2-sphere which we sample at the north-pole and at six points forming a regular hexagon along the equator. Note that any closed hemi-sphere contains at least three landmarks. It follows that the seven landmarks form an ε-sample of $\mathbb{M}$ for every $\varepsilon > \sqrt{2}$. Any tetrahedron spanned by the north-pole and three of the landmarks on the equator has a weak witness for each of its faces. Nevertheless, the tetrahedron does not have a strong witness on $\mathbb{M}$, which implies $\varepsilon_2 \leq \sqrt{2} = 1.414\ldots$.

6 Discussion

The main contribution of this paper is a proof that witness complexes as introduced in [de Silva and Carlsson 2004] provide a viable approach to reconstructing dimension 1 and dimension 2 submanifolds of Euclidean space. As recently established in [Oudot 2007], the situation is more complicated for submanifolds of dimension 3 and higher. Specifically, Oudot showed that for every $\varepsilon > 0$ there is a dimension 3 submanifold $\mathbb{M}$ of $\mathbb{R}^4$, an ε-sample L of $\mathbb{M}$, and a tetrahedron $\sigma \subseteq L$ such that every face of σ has a weak witness in $\mathbb{M}$ but σ has no strong witness in $\mathbb{M}$. It follows that the weak Delaunay constant for dimension 3 submanifolds vanishes. It is not difficult to extend this construction to submanifolds of dimension $k > 3$. Hence, $\varepsilon_k = 0$ for $k \geq 3$. The contrast to $\varepsilon_k > 0$ for $k = 1, 2$ as established in this paper highlights an interesting difference between at most two and at least three dimensions. The case of dimension 2 submanifolds in $\mathbb{R}^3$ is perhaps the most important one for applications. Here we prove $\frac{1}{\sqrt{5}} \leq \varepsilon_2 \leq \sqrt{2}$. Can we narrow the gap or perhaps even determine the weak Delaunay constant, at least for surfaces in $\mathbb{R}^3$? It would also be interesting to shed light on the dependence of the constant on the dimension of the ambient space. Clearly the constant cannot increase when that dimension goes up, that is, $\varepsilon_{2,d} \geq \varepsilon_{2,d'}$ whenever $d \leq d'$. Is this inequality strict? For example, is it true that $\varepsilon_{2,3} < \varepsilon_{2,4}$?

The counterexample by Oudot rests on the existence of slivers which are tetrahedra in the Delaunay triangulation in $\mathbb{R}^4$ with instable direction of their dual Voronoi edges. Building on work in [Cheng et al. 2000; Cheng et al. 2005], Boissonnat, Guibas and Oudot suggest to overcome this difficulty by assigning weights to the landmarks in such a way that the weighted Delaunay triangulation does not contain any slivers [Boissonnat et al. 2007]. To get a positive result, they use samples of dimension k submanifolds that are both sufficiently fine and sufficiently coarse. Beyond using weights for landmarks, they also propose to use weights for the witnesses as a means to counteract the imperfections in the reconstruction caused by the use of finite sets of witnesses in practice. This should be compared to the use of a tolerance parameter, β, in the construction of the family of witness complexes in [Attali et al. 2007].

References

AMENTA, N., AND BERN, M. 1999. Surface reconstruction by voronoi filtering. *Discrete Comput. Geom. 22*, 481–504.

AMENTA, N., ATTALI, D., AND DEVILLERS, O. 2007. Complexity of delaunay triangulation for points on lower-dimensional polyhedra. In *Proc. 18th Ann. ACM-SIAM Sympos. Discrete Alg.* to appear.

ATTALI, D., BOISSONNAT, J.-D., AND LIEUTIER, A. 2003. Complexity of the delaunay triangulation of points on surfaces: the smooth case. In *Proc. 19th Ann. Sympos. Comput. Geom.*, 201–210.

ATTALI, D., EDELSBRUNNER, H., HARER, J., AND MILEYKO, Y. 2007. Alpha-beta witness complexes. Manuscript, Duke Univ., Durham, North Carolina.

BERNARDINI, F., MITTLEMAN, J., RUSHMEIER, H., SILVA, C., AND TAUBIN, G. 1999. The ball-pivoting algorithm for surface reconstruction. *IEEE Trans. Visual. Comput. Graphics 5*, 349–359.

BOISSONNAT, J.-D., GUIBAS, L. J., AND OUDOT, S. Y. 2007. Manifold reconstruction in arbitrary dimensions using witness complexes. In *Proc. 23rd Ann. Sympos. Comput. Geom.* to appear.

CARR, J. C., BEATSON, R. K., CHERRIE, J. B., MITCHELL, T. J., FRIGHT, W. R., MCCALLUM, B. C., AND EVANS, T. R. 2001. Reconstruction and representation of 3d objects with radial basis functions. In *SIGGRAPH'01 Proceedings, Comput. Graphics*, 67–76.

CHAZAL, F., AND LIEUTIER, A. 2006. Topology guaranteeing manifold reconstruction using distance function to noisy data. In *Proc. 22nd Ann. Sympos. Comput. Geom.*, 112–118.

CHAZAL, F., COHEN-STEINER, D., AND LIEUTIER, A. 2006. A sampling theory for compact sets in euclidean space. In *Proc. 22nd Ann. Sympos. Comput. Geom.*, 319–336.

CHENG, S.-W., DEY, T. K., EDELSBRUNNER, H., FACELLO, M. A., AND TENG, S.-H. 2000. Sliver exudation. *J. Assoc. Comput. Mach. 47*, 883–904.

CHENG, S.-W., DEY, T. K., AND RAMOS, E. A. 2005. Manifold reconstruction from point samples. In *Proc. 16th Ann. ACM-SIAM Sympos. Discrete Alg.*, 1018–1027.

DE SILVA, V., AND CARLSSON, G. 2004. Topological estimation using witness complexes. In *Proc. Sympos. Point-Based Graphics*, 157–166.

DE SILVA, V. 2003. A weak definition of delaunay triangulation. Manuscript.

DO CARMO, M. P. 1976. *Differential Geometry of Curves and Surfaces*. Prentice Hall, Upper Saddle River, New Jersey.

EDELSBRUNNER, H., AND SHAH, N. R. 1997. Triangulating topological spaces. *Internat. J. Comput. Geom. Appl. 7*, 365–378.

EDELSBRUNNER, H. 2004. Surface reconstruction by wrapping finite point sets in space. In *Discrete and Computational Geometry — The Goodman-Pollack Festschrift*, B. Aronov, S. Basu, J. Pach, and M. Sharir, Eds. Springer-Verlag, Berlin, 379–404.

GIESEN, J., AND JOHN, M. 2002. Surface reconstruction based on a dynamical system. *Comput. Graphics Forum 21*, 363–371.

GUIBAS, L. J., AND OUDOT, S. Y. 2007. Reconstruction using witness complexes. In *Proc. 18th Ann. ACM-SIAM Sympos. Discrete Alg.* to appear.

HOPPE, H., DEROSE, T., DUCHAMP, T., MCDONALD, J., AND STUETZLE, W. 1992. Surface reconstruction from unorganized points. In *SIGGRAPH'92 Proceedings, Comput. Graphics*, vol. 26, 71–77.

MARTINETZ, T., AND SCHULTEN, K. 1994. Topology representing networks. *Neural Networks 7*, 507–522.

MILNOR, J. 1963. *Morse Theory*. Princeton Univ. Press, New Jersey.

NIYOGI, P., SMALE, S., AND WEINBERGER, S. to appear. Finding the homology of submanifolds with high confidence from random samples. *Discrete Comput. Geom.*.

OUDOT, S. Y. 2007. On the topology of the restricted delaunay triangulation and witness complexes in higher dimensions. Manuscript, Dept. Comput. Sci., Stanford Univ., California.

THIBERT, B. 2003. *Approximation géométrique d'une surface lisse. Applications en géologie structurale*. PhD thesis, Université Claude Bernard, Lyon, France.

ZHAO, H., OSHER, S., AND FEDKIW, R. 2001. Fast surface reconstruction using the level set method. In *Proc. 1st IEEE Workshop on Variational and Level Set Methods*, 194–202.

Appendix A

In this appendix, we present basic inequalities relating distances between points on a 2-manifold and angles between vectors the points define. Similar inequalities can be found in [Thibert 2003]. We introduce some notation, referring to [do Carmo 1976] for additional background. Let $\mathbb{M}$ be a smoothly embedded, compact 2-manifold in $\mathbb{R}^d$. For $x \in \mathbb{M}$, we write $\mathrm{T}\mathbb{M}_x$ for the tangent space of $\mathbb{M}$ at x and $\mathrm{N}\mathbb{M}_x = \mathrm{T}\mathbb{M}_x^\perp$ for its normal space. The tangent space has dimension 2 and the normal space has dimension $d-2$. If $\gamma : \mathbb{R} \to \mathbb{M}$ is a smooth, arc-length parametrized curve, the unit tangent and normal vectors at $x = \gamma(0)$ are $T_\gamma(x) = \dot\gamma(0)$ and $N_\gamma(x) = \ddot\gamma(0)/\|\ddot\gamma(0)\|$. The curvature of γ at x is $\kappa_\gamma(x) = \|\ddot\gamma(0)\|$. If γ is a geodesic then $\kappa_\gamma(x)$ is as small as it can be for a curve that passes through x in the direction $u = T_\gamma(x) \in \mathrm{T}\mathbb{M}_x$. This is also the curvature of the normal section obtained by intersecting $\mathbb{M}$ with the $(d-1)$-dimensional plane containing x, $x+u$, and $x + \mathrm{N}\mathbb{M}_x$. In this case, we call $\kappa_u(x) = \kappa_\gamma(x)$ the sectional curvature of $\mathbb{M}$ at x in the tangent direction u. We write $\kappa(x) = \max_u \kappa_u(x)$ for the maximum (absolute) sectional curvature at x, and $\kappa = \max_x \kappa(x)$ for the maximum (absolute) sectional curvature anywhere on $\mathbb{M}$.

Distance. Let $d(x,y)$ be the length of the shortest path connecting the points x and y on $\mathbb{M}$. If $l = d(x,y)$ then there is an arc-length parametrization of a geodesic, $\gamma : [0,l] \to \mathbb{M}$, with $x = \gamma(0)$ and $y = \gamma(l)$. Because γ is unit-speed, its derivative is a map to the unit sphere, $\dot\gamma : [0,l] \to \mathbb{S}^{d-1}$. For every $0 \le s \le l$, the angle between the vectors $\dot\gamma(0)$ and $\dot\gamma(s)$ is bounded from above by the length of the path connecting them on the unit sphere: $\angle\dot\gamma(0)\dot\gamma(s) \le \int_0^s \|\ddot\gamma(t)\|\,\mathrm{d}t \le \kappa s$. The second inequality follows from $\|\ddot\gamma(t)\| \le \kappa$ for all t since γ is a geodesic parametrized by arc-length. We use this fact to bound the Euclidean distance between x and y in terms of the geodesic distance.

PROPERTY I. $\|x-y\| \ge \frac{2}{\kappa}\sin(\frac{\kappa}{2}d(x,y))$, provided the geodesic distance between x and y is $d(x,y) \le \frac{\pi}{\kappa}$.

PROOF. Let $T = \dot\gamma(\frac{l}{2})$ be the unit tangent vector at the halfway point. The length of γ in the direction T is

$$\begin{aligned}\langle y-x, T\rangle &= \int_{s=0}^{l} \langle\dot\gamma(s), T\rangle\,\mathrm{d}s \\ &\ge 2\int_{s=0}^{l/2} \cos(\kappa s)\,\mathrm{d}s\end{aligned}$$

because the angle between T and $\dot\gamma(s)$ is at most $\kappa|\frac{l}{2} - s|$. The right hand side of the inequality evaluates to $\frac{2}{\kappa}\sin(\frac{\kappa}{2}l)$. The length of γ in the direction $y-x$ can only be larger, which implies the claimed inequality. ∎

The bound in Property I is tight, with equality realized by an arc of a circle with radius $\frac{1}{\kappa}$.

Angle between tangent spaces. Given two tangent spaces, we define their *angle* as the max-min angle between any two of their vectors,

$$\angle \mathrm{T}\mathbb{M}_x \mathrm{T}\mathbb{M}_y \quad = \quad \max_{u \in \mathrm{T}\mathbb{M}_x} \min_{v \in \mathrm{T}\mathbb{M}_y} \angle uv.$$

Note that it satisfies the triangle inequality, $\angle \mathrm{T}\mathbb{M}_x \mathrm{T}\mathbb{M}_z \leq \angle \mathrm{T}\mathbb{M}_x \mathrm{T}\mathbb{M}_y + \angle \mathrm{T}\mathbb{M}_y \mathrm{T}\mathbb{M}_z$, for all points x, y, z. We bound the angle between two tangent spaces in terms of the geodesic distance between their base points. Instrumental to this proof is the notion of parallel transport from $\mathrm{T}\mathbb{M}_x$ to $\mathrm{T}\mathbb{M}_y$ along a smooth curve $\gamma : [0,1] \to \mathbb{M}$ connecting $x = \gamma(0)$ with $y = \gamma(1)$. To explain this concept, consider a vector field $u : [0,1] \to \mathbb{R}^d$ that assigns to each $s \in [0,1]$ a tangent vector $u(s) \in \mathrm{T}\mathbb{M}_{\gamma(s)}$. This vector field is *parallel* if its covariant derivative vanishes, that is, the orthogonal projection of $\frac{\mathrm{d}u}{\mathrm{d}s}(s)$ onto $\mathrm{T}\mathbb{M}_{\gamma(s)}$ is zero for every $s \in [0,1]$. Given an initial tangent vector, $u(0) \in \mathrm{T}\mathbb{M}_x$, the parallel vector field exists and is unique [do Carmo 1976, Chapter 4]. The vectors $u(t)$ are then called the *parallel transport* of $u(0)$ along γ. As it turns out, the implied map from $\mathrm{T}\mathbb{M}_x$ to $\mathrm{T}\mathbb{M}_y$ preserves scalar products and is therefore an isometry. This map also implies an upper bound on the angle between the two tangent spaces.

PROPERTY II. $\angle \mathrm{T}\mathbb{M}_x \mathrm{T}\mathbb{M}_y \leq \kappa d(x,y)$.

PROOF. As before, we set $l = d(x,y)$ and let $\gamma : [0,l] \to \mathbb{M}$ be an arc-length parametrization of a geodesic with $x = \gamma(0)$ and $y = \gamma(l)$. Given a unit vector $u(0) \in \mathrm{T}\mathbb{M}_x$, we let $u(t)$ be its parallel transport along $\gamma(t)$. Because the $u(t)$ are unit vectors, u is a map to the unit sphere, $u : [0,l] \to \mathbb{S}^{d-1}$, whose derivative is bounded by the maximum sectional curvature, $\|\dot{u}(t)\| \leq \kappa$. The angle between the unit vectors $u(0)$ and $u(l)$ is the length of the great-circle arc connecting them on $\mathbb{S}^{d-1}$, which is bounded by the length of the path on the sphere, $\angle u(0)u(l) \leq \int_0^l \|\dot{u}(t)\| \,\mathrm{d}t \leq \kappa l$. The claim follows by choosing $u(0) \in \mathrm{T}\mathbb{M}_x$ and $v \in \mathrm{T}\mathbb{M}_y$ such that $\angle \mathrm{T}\mathbb{M}_x \mathrm{T}\mathbb{M}_y = \angle u(0)v \leq \angle u(0)u(l)$. □

Similar to Property I, the bound in Property II is tight, with equality realized by an arc of a circle with radius $\frac{1}{\kappa}$. To make this arc into a geodesic, we may place it on a sphere with the same radius.

Angle to tangent space. Consistent with the notion of angle between two tangent spaces, we define the *angle* between $v = y - x$ and the tangent space at x equal to $\angle v\mathrm{T}\mathbb{M}_x = \min_{u \in \mathrm{T}\mathbb{M}_x} \angle vu$. Using a proof similar to that of Property I, we bound this angle in terms of the maximum sectional curvature and the geodesic distance between the points.

PROPERTY III. $\angle v\mathrm{T}\mathbb{M}_x \leq \frac{\kappa}{2} d(x,y)$, provided the geodesic distance between the two points is $d(x,y) \leq \frac{\pi}{2\kappa}$.

PROOF. As before, we set $l = d(x,y)$ and let $\gamma : [0,l] \to \mathbb{M}$ be an arc-length parametrization of a geodesic with $\gamma(0) = x$ and $\gamma(l) = y$. Let $T_0 = T_\gamma(0)$ and $N_0 = N_\gamma(0)$ be the tangent and normal vectors at x. The orthogonal projection of the vector $v = y - x$ onto the line spanned by N_0 and passing through x has length

$$\begin{aligned} |\langle v, N_0 \rangle| \quad &\leq \quad \int_{s=0}^{l} |\langle \dot{\gamma}(s), N_0 \rangle| \,\mathrm{d}s \\ &\leq \quad \int_{s=0}^{l} \sin(\kappa s) \,\mathrm{d}s, \end{aligned}$$

with the right hand side evaluating to $\frac{1}{\kappa}(1 - \cos(\kappa l))$. Similarly, the orthogonal projection of v onto the line spanned by T_0 and passing through x is

$$\begin{aligned} \langle v, T_0 \rangle \quad &= \quad \int_{s=0}^{l} \langle \dot{\gamma}(s), T_0 \rangle \,\mathrm{d}s \\ &\geq \quad \int_{s=0}^{l} \cos(\kappa s) \,\mathrm{d}s, \end{aligned}$$

with the right hand side evaluating to $\frac{1}{\kappa} \sin(\kappa l)$. The two inequalities are equalities if γ is an arc of a circle with center $x \pm \frac{1}{\kappa} N_0$ and radius $\frac{1}{\kappa}$ contained in the 2-plane passing through x, $x + T_0$, and $x + N_0$. In this special configuration, $\angle v T_0 = \arctan \frac{1 - \cos(\kappa l)}{\sin(\kappa l)} = \frac{\kappa l}{2}$. The claimed inequality follows. □

Similar to Properties I and II, the bound in Property III is tight, with equality realized by an arc of a circle whose radius is $\frac{1}{\kappa}$.