How Many Delaunay Cells Does a Segment Cross for Random Sites on a Manifold?

Pedro M. M. de Castro
Centro de Informática, Universidade Federal de Pernambuco
Recife, Pernambuco, Brazil
pmmc@cin.ufpe.br

November 23, 2026

hexagon:2611.00964v1

Abstract

When sites are sampled on the boundary of a solid, a straight walk through their Delaunay triangulation can traverse empty interior space far from every site. Does the spacing of sites on the boundary still govern how many cells the walk visits? Fix an ambient dimension $d \geq 2$ and a compact smooth manifold of intrinsic dimension $1 \leq \delta < d$ without boundary, and consider its smooth embeddings in $\mathbb{R}^d$ whose images lie in no affine hyperplane. We prove that the intrinsic spacing bound holds for a generic embedding, meaning a countable intersection of open dense sets in the C^∞ topology. Draw n sites independently and uniformly (or from a density bounded above and away from zero) on the embedded manifold. For almost every pair of endpoints inside its convex hull, the segment s between them meets $O(1 + |s|n^{1/\delta})$ Delaunay cells in expectation, where $|s|$ is its length. The implied constant may depend on where s lies, and stays uniform on compact subsets of an open region whose complement has zero $(d-1)$-dimensional Hausdorff measure. For volume samples in a convex body of dimension $d \geq 1$, we prove the expected upper bound $1 + C_A|s|n^{1/d}$ whenever s stays at distance at least $An^{-1/d}$ from the boundary, removing the logarithmic factor previously known in three dimensions.

By contrast, intrinsic dimension alone cannot guarantee this rate for every smooth shape. On an elongated spheroid, an empty ball can touch along a full circle. Under uniform surface sampling, the expected count stays $O(1)$ on a fixed segment inside one contact disk, yet grows linearly in n for a fixed interior segment of the same length with nonzero axial displacement. Even when all contacts are isolated, a curve in $\mathbb{R}^3$ can exceed the linear scale $O(n)$. We build both a smooth closed curve and a pair of disjoint real-analytic closed curves in $\mathbb{R}^3$ along which a moving empty sphere keeps three isolated contacts, one of fourth order, over a fixed interval, so that many crossed tetrahedra share one edge. Under uniform arclength sampling on the smooth curve, or on the union of the two analytic curves with weights proportional to length, with probability tending to one, at least $cn^{4/3}$ Delaunay tetrahedra simultaneously meet every segment joining two fixed interior balls. Moreover, for each $0 < \varepsilon < 1/3$, approximating this geometry on a shrinking window gives a connected real-analytic closed curve in $\mathbb{R}^3$ with simultaneous lower bound $\Omega(n^{4/3-\varepsilon})$.

AI-use disclosure. OpenAI's ChatGPT and Codex, and Anthropic's Claude, were used as research tools to assist the author with literature retrieval, mathematical exploration, symbolic and numerical checks, illustrations and draft preparation and revision. Antigravity assisted with editorial review. The author formulated the research questions and mathematical framework, developed and refined the results and arguments with this assistance, and critically reviewed the manuscript throughout. Responsibility for the mathematical statements and the final manuscript rests entirely with the author.

Keywords. Delaunay triangulation; point location; random manifold samples; segment crossings; contact geometry.

1 Introduction

A classic way to locate a query point in a Delaunay triangulation is the *straight walk*. One starts at a known cell and traces the line segment s to the query point, stepping into an adjacent cell whenever s crosses a shared facet. In d dimensions, each cell of a generic triangulation is a d-simplex (a triangle in the plane, or a tetrahedron in $\mathbb{R}^3$) whose open *circumball* contains none of the sampled points, called *sites*. Because the walk advances strictly from one simplex to the next along s, the number of Delaunay cells whose interiors meet s directly measures its work. How many cells does such a walk visit?

Figure 1 shows why this question turns subtle as soon as sites sit on a surface rather than filling a solid. In three dimensions, a small ball of radius r inside the solid ellipsoid of Fig. 1(a) has volume proportional to r^3. By contrast, a small patch of radius r on its boundary surface in Fig. 1(b) has area proportional to r^2. To keep the expected number of sites in such a neighbourhood roughly constant as the sample size n grows, r must scale as $n^{-1/3}$ in the solid and as $n^{-1/2}$ on the surface. Multiplying n by 64 therefore cuts the volume spacing scale by four, while cutting the surface spacing scale by eight. Yet in both Figs. 1(a) and (b), the segment $s = [a, b]$ runs deep inside the ellipsoid, far from every boundary site. When all vertices lie on the curved outer shell, why should the spacing of sites *on the boundary* still govern how many tetrahedra cross the empty interior?

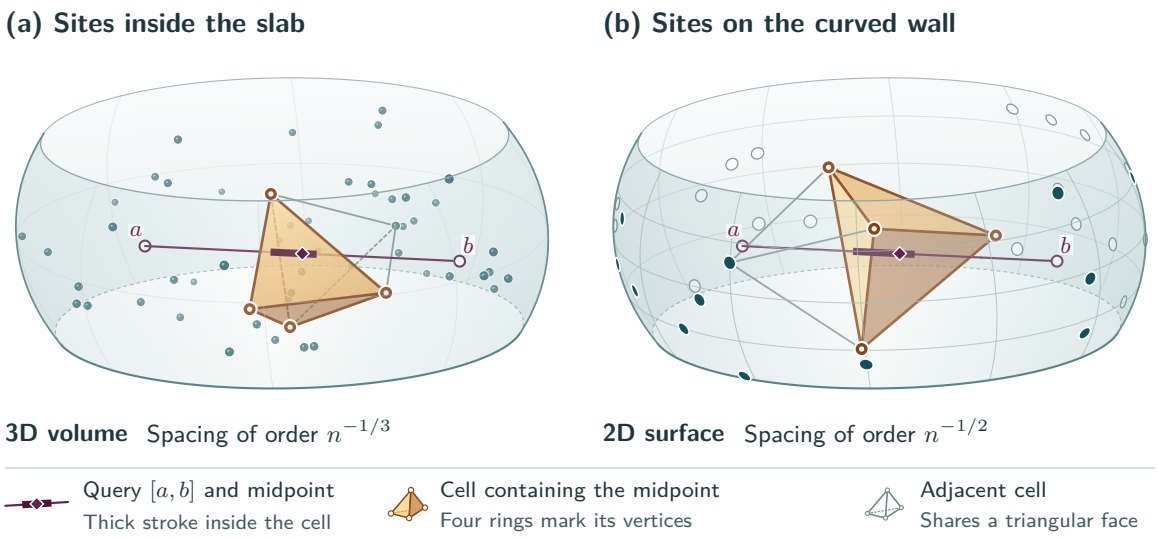


Figure 1: Two parallel cuts expose the same slab of a triaxial ellipsoid with a unique shortest semiaxis. In (a), sites in the solid appear at different depths. In (b), all sites, including the four vertices of the highlighted tetrahedron, lie on the curved wall. Dark blue dots mark sites on the near side, and open circles mark sites on the far side. Four brown rings mark the vertices of the tetrahedron containing the midpoint (diamond) of $s = [a, b]$, and the thick stroke marks the portion of s inside it; faint edges complete a neighbouring tetrahedron across a shared face.

De Castro and Devillers raised this question for sites distributed evenly on a bounded submanifold $V \subset \mathbb{R}^d$ of *intrinsic dimension* δ [11, Conjecture 2, p. 435]. Here δ is the dimension of the sampled manifold itself (so $\delta = 1$ for a curve and $\delta = 2$ for a surface), rather than the ambient dimension d in which the triangulation lives. For a segment s of Euclidean length $|s|$, they conjectured an upper bound of the form $1 + C|s|n^{1/\delta}$ on the expected number of visited cells. They also tested a randomized walk that chooses the next cell at random among admissible neighbours, observing experimental exponents of 0.52 on a triaxial ellipsoid and 0.51 on a closed cylinder [11, Sections 2.1 and 7.1]. Here we analyse the deterministic straight walk along the segment s itself.

To formalize the model, fix a segment $s \subset \mathbb{R}^d$ before drawing the sites, and let $N_n(s)$ count the full-dimensional Delaunay simplices whose interiors meet s. Its expectation $\mathbb{E}N_n(s)$ is taken over the n random sites. The Euclidean metric on $\mathbb{R}^d$ induces a natural δ-dimensional volume on V, called its *intrinsic volume* (arclength when $\delta = 1$, surface area when $\delta = 2$). Dividing the intrinsic volume of any measurable subset $E \subset V$ by the total intrinsic volume of V yields its *normalized intrinsic volume* $v \in [0, 1]$. Under uniform sampling, a single site falls in E with probability v. More generally, our upper bounds allow any independent sampling law μ on V whose density relative to normalized intrinsic volume stays bounded above and away from zero. That is, there exist fixed constants $0 < m \leq 1 \leq M < \infty$, independent of E and n, such that every measurable subset $E \subset V$ satisfies

$$mv \leq \mu(E) \leq Mv. \tag{1.1}$$

For example, if an arc accounts for one tenth of a closed curve's length, each draw lands on that arc with probability between $m/10$ and $M/10$. Uniform sampling corresponds to $m = M = 1$.

Our main positive result, Theorem 1.1, fixes an abstract compact smooth δ-manifold X without boundary and varies how X sits in $\mathbb{R}^d$. Specifically, we consider smooth embeddings of X into $\mathbb{R}^d$ (injective smooth maps whose derivative is injective everywhere) whose images have *full affine span*, meaning they lie in no affine hyperplane. We equip this space with the C^∞ topology, where closeness requires uniform convergence of the maps and all their derivatives on compact coordinate charts. A class of embeddings is *residual* if it contains a countable intersection of open dense subsets, and a property holds for a *generic embedding* if it holds on a residual class. Intuitively, this is the smooth analogue of general position for finite point sets: just as tiny nudges remove collinear triples or cocircular quadruples in the plane, an arbitrarily small smooth perturbation of V eliminates accidental sphere tangencies. Once an embedding with image V is fixed, the phrase *almost every pair of endpoints* refers to Lebesgue measure on $\mathbb{R}^d \times \mathbb{R}^d$.

How does generic position enter the geometry of the walk? As Section 2 explains, the cells met by s are controlled by empty spheres tangent to V. We describe those tangent spheres by systems of equations involving the contact points on V, the sphere centre and radius, and auxiliary parameters for a supporting hyperplane or a direction of zero second derivative. Roughly speaking, an embedding is generic in our sense when every such defining system has full row rank at every solution. Varying the contact points and geometric parameters together then allows independent first-order changes in the equations. That transversality prevents tangent spheres from developing accidental degeneracies or extra contacts beyond their expected codimension, reducing the exceptional query sets to the five low-dimensional cases catalogued in Subsections A.2 and A.4.

Theorem 1.1. *Fix integers $d \geq 2$ and $1 \leq \delta \leq d - 1$, the ambient and intrinsic dimensions. Let X be a compact smooth δ-manifold without boundary admitting a smooth embedding in $\mathbb{R}^d$ whose image lies in no affine hyperplane. In the C^∞ topology, a residual class of these embeddings has the following property.*

Fix such an embedding, let V be its image and write $\operatorname{int} K$ *for the interior of its convex hull K. Draw n sites independently on V from a common probability law μ satisfying* (1.1) *for every measurable subset of V. For Lebesgue-almost every ordered pair $(a, b) \in (\operatorname{int} K)^2$, there are constants $C_s > 0$ and $n_s \in \mathbb{N}$ such that the segment $s = [a, b]$ satisfies*

$$\mathbb{E}N_n(s) \leq C_s(1 + |s|n^{1/\delta}) \qquad (n \geq n_s).$$

The embedding determines an open set $G \subset \operatorname{int} K$. For each compact $Q \subset G$, the bound holds for every segment $s \subset Q$ fixed before sampling, with common constants $C_Q > 0$ and $n_Q \in \mathbb{N}$ depending only on Q, the embedding and the bounds m, M in (1.1). *For every $\varepsilon > 0$, the*

complement $(\operatorname{int} K) \setminus G$ *can be covered by countably many balls of radii* r_i *with* $\sum_i r_i^{d-1} < \varepsilon$, *meaning it has zero* $(d-1)$*-dimensional Hausdorff measure.*

Appendix A proves Theorem 1.1, completing the segment estimate on page 43. Notice why $(\operatorname{int} K) \setminus G$ must have zero $(d-1)$-dimensional Hausdorff measure rather than merely zero d-dimensional volume. A zero-volume hypersurface could still split an open ball in two and block an open family of segments. By contrast, a set of zero $(d-1)$-dimensional Hausdorff measure projects to a null set on every hyperplane, so almost every line in any fixed direction misses it, and Fubini's theorem turns that into almost every pair $(a, b) \in (\operatorname{int} K)^2$ (page 35). For any two-dimensional manifold in this residual class, the theorem gives $\mathbb{E}N_n(s) = O(1 + |s|\sqrt{n})$ in every ambient dimension $d \geq 3$. Unlike [11, Conjecture 2], which asks for one uniform coefficient of $|s|n^{1/\delta}$ across the entire hull and an additive term of 1, Theorem 1.1 allows C_s and n_s to depend on where the segment lies inside K.

Even for a concrete shape outside this abstract residual class, one can often check the contact hypotheses directly. Return to the triaxial ellipsoid of Fig. 1, which has a unique shortest semiaxis. For the same interior segment, Proposition 2.1 bounds the expected number of tetrahedra visited by $1 + C|s|n^{1/3}$ when sites are uniform inside the solid, whereas Proposition 4.4 gives $C'(1 + |s|\sqrt{n})$ when sites are uniform on the boundary surface. Moreover, for any fixed interior segment parallel to the shortest axis, the surface expectation stays $O(1)$ as n grows.

How surface sites control an interior walk. To prove an upper bound of order $1+|s|n^{1/\delta}$, we split any segment s of positive length into at most $1 + |s|n^{1/\delta}$ short pieces of length at most $n^{-1/\delta}$. We then bound the expected number of simplices meeting each piece by a constant independent of n. Because a piece may sit deep inside K while all sites lie on V, we link interior queries $q \in s$ to the manifold through empty balls. Given $q \in \operatorname{int} K$, consider a ball whose open interior contains no point of V. The points where its boundary sphere touches V are its *contacts*, and we say that the ball *supports* q if q lies in the convex hull of those contacts.

For a ball B of centre c and radius R, the *power* of a point $p \in \mathbb{R}^d$ is $|p - c|^2 - R^2$. It vanishes at every contact of B and stays nonnegative on V. Call a contact *ordinary* if the restriction of the power to V has a positive definite Hessian there, so moving a small distance t along V raises the power at least quadratically in t. We say that q is *regular* when two conditions hold: every contact of every ball supporting q is ordinary, and the tangent spaces at any minimal representing set of contacts together with the chords joining them span $\mathbb{R}^d$. At a regular query, sampling sites inside small patches around the contacts encloses a neighbourhood of q of radius proportional to t. Comparing the supporting sphere with the circumsphere of any cell meeting that neighbourhood forces all vertices of the cell into small $O(t)$-neighbourhoods of the contacts. Lemma 3.1 then bounds the local expected count at scale $t = n^{-1/\delta}$. Notably, for a fixed segment s contained in a compact set of regular queries, if one ball supports every point of s, the same witness patches serve the whole segment at once, keeping $\mathbb{E}N_n(s)$ bounded as n grows.

Along a general segment, however, the supporting contacts need not stay fixed. The walk can cross a configuration where the tangent directions and chords of the representing contacts span only a hyperplane $W \subset \mathbb{R}^d$. For a generic embedding, the dimension counts in Subsection A.2 show that almost every segment avoids all such deficient sets unless $d = j(\delta+1)$, where j is the number of contacts in a minimal representation of q. In $\mathbb{R}^3$ with a surface $(\delta = 2)$, the sole divisible case is $j = 1$. That simply describes a segment piercing a nonconvex surface V inside its convex hull: at the crossing point $q \in V$, the radius-zero ball centred at q bridges the largest empty tangent balls on the two sides of T_qV. At a typical point of any codimension-one interface, tangency and equidistance at the j base contacts leave one free

coordinate for the sphere centre along the line normal to W. Emptiness confines the centre to a closed interval whose two endpoints each gain one extra contact, located on opposite sides of the hyperplane $q + W$. Sites sampled near those two extra contacts supply the missing normal direction on their respective sides, allowing the local bound to cross the interface cleanly (Section 3 and Fig. 6).

When contact geometry breaks the bound. Why are the generic hypotheses of Theorem 1.1 genuinely needed? Two explicit constructions show how non-generic contacts can break the $n^{1/\delta}$ scale in opposite ways. The first occurs on a prolate spheroid, an ellipsoid of revolution stretched along its symmetry axis (Section 4). There an empty ball can touch the surface along an entire circular latitude. With uniform surface sites, a fixed segment inside the disk spanned by one contact circle meets only $O(1)$ cells in expectation. By contrast, a fixed interior segment of the same length with nonzero displacement along the long axis crosses a continuum of contact disks and visits $\Theta(n)$ cells in expectation (Corollary 4.3). Locating cell vertices inside thin axial belts explains this jump and extends the $\Theta(n)$ obstruction to smooth analytic products in every intrinsic dimension $2 \le \delta \le d-1$.

A second obstruction arises on curves in $\mathbb{R}^3$ ($d = 3$, $\delta = 1$), where the spacing bound already permits linear growth $O(n)$ and even isolated contacts can force superlinear crossings (Section 5). The mechanism uses a one-parameter family of empty spheres that touch the curve at only three isolated points. At one front contact, the sphere power grows as the fourth power of arclength, while at the two rear contacts it grows quadratically. Each suitably isolated pair of sites near the rear contacts forms a single edge shared by order $n^{1/3}$ crossed tetrahedra whose opposite edges are consecutive sites near the flatter fourth-order contact. Both a single C^∞ closed curve and a pair of disjoint real-analytic closed curves maintain these contact orders over a fixed interval. Under uniform arclength sampling on the smooth curve, or on the union of the two real-analytic curves with component probabilities proportional to their lengths, on an event of probability at least $1 - \exp(-c_1 n^{1/3})$ at least $cn^{4/3}$ Delaunay tetrahedra simultaneously meet *every* segment joining two fixed interior balls. Moreover, for each $0 < \varepsilon < 1/3$, approximating this geometry over a shrinking interval yields a fixed connected real-analytic closed curve in $\mathbb{R}^3$ whose corresponding simultaneous lower bound is $\Omega(n^{4/3-\varepsilon})$; whether a single connected real-analytic closed curve can attain the exact exponent $4/3$ remains open. Section 6 complements these theorems by measuring how intrinsic dimension, segment length and contact geometry behave in finite samples.

Earlier estimates for volume samples. Even when sites fill a full-dimensional convex body rather than a lower-dimensional manifold, two subtleties shape the literature: how close the segment may approach the boundary, and whether the segment is fixed before or chosen after sampling. In the plane, Bose and Devroye [6, Theorem 1] proved that for all sufficiently large n, any fixed segment of length ℓ inside a convex body crosses at most $C\ell\sqrt{n}$ Voronoi edges in expectation when the area density stays bounded above and away from zero. For uniform sites in a planar convex body of area one, Devroye, Mücke and Zhu [16, Section 2, p. 479] stated the bound $c_3 + c_4\ell\sqrt{n}$ for Delaunay triangles at boundary distance at least $3\sqrt{\log n/n}$. Devroye, Lemaire and Moreau [15, Lemma 4] proved the analogous bound for crossed Delaunay edges in the unit square at the same boundary margin. In sharp contrast, if one maximizes over full lines chosen after observing n uniform sites in a planar convex body rather than fixing a segment beforehand, the expected maximum number of crossed Delaunay edges is $\Theta(\sqrt{n})$ [7, Theorem 1].

In three dimensions, Mücke, Saias and Zhu [27, Theorem 2 and its proof, pp. 66 and 77] studied independent sites in a bounded convex body K with volume density bounded above and away from zero. For the walk cost along a segment at boundary distance $c(\log n/n)^{1/3}$,

they proved a crossing bound carrying an extra factor $\log n / \log\log n$, and asked whether this logarithmic factor could be removed and how walks behave in higher dimensions. The $\ell n^{1/d}$ length dependence up to the boundary also appeared in de Castro's thesis [10, p. 100]. In arbitrary dimension d, Hemsley [20, Theorem 3.7 and Remark 3.8] stated a binomial $\ell n^{1/d}$ bound, carried out the proof for a homogeneous Poisson point process on $\mathbb{R}^d$ (assuming $d \geq 2$ at the end), and referred to his earlier de-Poissonization section for fixed n. Here, Proposition 2.1 gives a direct fixed-n proof for any bounded convex body $K \subset \mathbb{R}^d$ of arbitrary dimension $d \geq 1$ and any density satisfying (1.1): for each fixed $A > 0$, every segment of length ℓ fixed before sampling at distance at least $An^{-1/d}$ from ∂K crosses at most $1 + C_A \ell n^{1/d}$ cells in expectation, removing the logarithmic factor and shrinking the boundary margin of [27] in three dimensions.

For an infinite homogeneous Poisson process in $\mathbb{R}^d$ of *intensity* λ (meaning λv sites on average in every region of volume v), translation and rotation invariance make the walk exact in expectation. Write c_d for the mean number of facets crossed per unit length at intensity one. Integral-geometric formulas relating crossings to projected cell sizes [31, proof of Theorem 4.1.2] and moment bounds [22, Section 2, Lemma 1] show that c_d is finite, so rescaling to intensity λ yields $1 + c_d \ell \lambda^{1/d}$ expected cells. In the plane this constant is $c_2 = 64/(3\pi^2)$ [31, Theorem 10.2.9], and Chenavier and Devillers [9, Proposition 6 of the preprint] obtained the same leading coefficient, up to an $O(1)$ remainder, between two inserted sites at unit distance. When the intensity is merely bounded above and away from zero relative to Lebesgue measure, Pimentel and Rossignol [28, Section 4.3, Lemma 4.11] bounded the expected maximum number of cells met by any line in a cube of side $L \geq 1$ by $O(L)$.

From the whole triangulation to a single walk. A straight walk typically visits only a tiny fraction of all Delaunay cells, yet both quantities are governed by empty-sphere contacts. Most earlier results on random surface samples studied the *total* complexity of the triangulation. For C^5 surfaces in $\mathbb{R}^3$ satisfying their genericity hypotheses, Attali, Boissonnat and Lieutier [4, Theorem 22] bounded total Delaunay complexity under deterministic (ε, κ)-sampling, and Devillers and Duménil [12, Theorem 3.7] established expected total complexity bounds for Poisson samples on a generic closed smooth surface of unit area. Duménil [17, Theorems 13.5 and 18.6] used contact geometry to bound expected total complexity on an oblate spheroid and on closed oriented surfaces under his genericity and contact hypotheses, distinguishing sphere centres on the two sides of a nonconvex surface's tangent plane [17, Section 15.1.2, p. 166]. On the finite lateral surface of a right circular cylinder (distinct from the closed cylinder tested experimentally in [11]), Devillers, Erickson and Goaoc [13, Section 4] proved an expected total complexity of $\Theta(n \log n)$ for n uniform sites. Duménil [17, Theorem 12.1 and Conjecture 12.9] obtained the matching Poisson bound for the expected total complexity on a right circular cylinder and conjectured the same $\Theta(\lambda \log \lambda)$ expected total complexity under Poisson sampling for any closed piecewise C^2 surface in $\mathbb{R}^3$ that is either a canal surface with no spherical part or a cylinder not degenerated into a plane, while Amenta, Attali and Devillers [1] proved tight worst-case bounds on polyhedral surfaces under deterministic sampling. Moving from total complexity in $\mathbb{R}^3$ to an individual walk in arbitrary dimensions, Theorem 1.1 controls the expected number of full-dimensional Delaunay simplices met by a fixed segment via intrinsic dimension δ, for a residual class of smooth embeddings in any dimensions $1 \leq \delta < d$, for almost every pair of interior endpoints, and with uniform constants on compact subsets of the good region G.

To isolate only the cells visited by a walk, we combine ideas from distance-function criticality and randomized convex hulls. Arnal, Cohen-Steiner and Divol [3, Lemmas 3.1, 4.1 and 4.8] analysed how nearest points on a generic submanifold behave near an ordinary contact and near a critical point of the distance function (a centre in the convex hull of its nearest points).

For our walk, the query q itself must lie in the convex hull of the contacts, which creates a one-dimensional interval of supporting centres at codimension-one transitions. Once the supporting contacts are identified, we count nearby cells using the witness and collector method of Devillers, Glisse, Goaoc and Thomasse [14, Section 1.2]. Placing sites inside a bounded collection of witness patches forces all vertices of every cell reaching q into a small collector region on V.

Examples with many cells. Our lower bounds also sit alongside known constructions of unexpectedly large Delaunay triangulations or long walks. On the Clifford torus (the product of two circles in orthogonal planes of $\mathbb{R}^4$), Magazinov [24, Theorems 2.1–2.3] estimated the face numbers of the convex hull of a Poisson sample. Replacing one circle by an ellipse in Section 4 breaks cosphericity and confines cells to narrow axial belts. On curves and surfaces in $\mathbb{R}^3$, Erickson [18, Corollary 2.2 and Section 3.5] proved that every pair of points on one open turn ($-\pi < t < \pi$) of a circular helix forms a Delaunay edge, and also constructed connected smooth surfaces chosen as a function of n whose total Delaunay complexity is nearly quadratic under independent sampling. A single polytope can also pair huge total complexity with small planar sections: a cyclic four-polytope on n vertices has quadratically many facets, yet any generic plane cross-section has at most $3n$ sides [2, Theorem 5.14 and the ensuing remark]. For finite point sets in $\mathbb{R}^3$, Shewchuk [32] constructed configurations where a line crosses $\Omega(n^2)$ tetrahedra. By contrast, Sections 4 and 5 fix a compact smooth manifold, a uniform sampling law, and a query segment or pair of endpoint balls independently of n: on the prolate spheroid and its higher-dimensional products, fixed axial segments reach $\Theta(n)$ expected crossings (exceeding $n^{1/\delta}$ for $\delta \geq 2$). On our smooth closed curve and two-component real-analytic link in $\mathbb{R}^3$ (sampled uniformly by arclength on the smooth curve or on the union of the two components with probabilities proportional to their lengths), with probability tending to one, at least $cn^{4/3}$ Delaunay tetrahedra simultaneously meet every segment joining two fixed interior balls; and for each $0 < \varepsilon < 1/3$, our connected real-analytic closed curve attains the corresponding simultaneous $\Omega(n^{4/3-\varepsilon})$ lower bound on a shrinking phase window.

Organization of the paper. Section 2 summarizes the recurring notation, proves the solid-body bound (Proposition 2.1), and introduces the paraboloid lift connecting interior queries to supporting spheres on V. Section 3 develops the witness-and-collector lemma, states the fixed-embedding bound (Theorem 3.2), and compares regular queries with codimension-one contact switches on explicit sphere embeddings. Sections 4 and 5 present our two counterexample families, on prolate spheroids/products and on space curves with a fourth-order contact. Section 6 reports finite-sample experiments on dimension, four-site fans and segment length, and Section 7 synthesizes the broader picture and open problems. Readers seeking the main geometric ideas first can read Sections 2–7 straight through before turning to the technical proofs in Appendices A–D.

2 From a Delaunay walk to its supporting spheres

Before analysing walks on manifolds, we set up our common counting framework for both sampling models and establish the solid-volume bound as a warm-up. When sites fill a solid convex body, empty cones inside the body bound the crossed facets directly along the segment. When sites lie on a lower-dimensional manifold, a segment deep inside the convex hull can stay far from every site. We therefore convert empty spheres into supporting hyperplanes of a paraboloid lift in $\mathbb{R}^{d+1}$, linking interior query points back to their contacts on the manifold.

Throughout the paper, $\operatorname{conv} V$ denotes the convex hull of a set V, and $\operatorname{int} K$ denotes the interior of K in the ambient Euclidean space $\mathbb{R}^d$. In the *volume model*, the ambient dimension is

an integer $d \geq 1$ and the sites are drawn inside a compact convex body $K \subset \mathbb{R}^d$ with nonempty interior. In the *manifold model*, the ambient dimension is an integer $d \geq 2$, the sites lie on an embedded manifold $V \subset \mathbb{R}^d$ of intrinsic dimension $1 \leq \delta \leq d-1$, and $K = \operatorname{conv} V$ again has nonempty interior. In both models, we fix the query segment $s \subset \operatorname{int} K$ before drawing n independent sites from a law satisfying (1.1). That reference measure is normalized Lebesgue volume on K in the volume model and normalized intrinsic volume on V in the manifold model. For convenience, the table below gathers the recurring symbols used throughout the paper.

Symbol	Meaning
$d,\ \delta$	Ambient dim. ($d \geq 1$ volume; $d \geq 2$ manifold) and intrinsic dim. ($1 \leq \delta < d$)
$V,\ K,\ \operatorname{int} K$	Embedded manifold, convex hull $K = \operatorname{conv} V$ (or body K), interior $\operatorname{int} K$
$s,\ \lvert s\rvert$	Fixed query segment in $\operatorname{int} K$ and its Euclidean length
$N_n(s),\ \Xi_n(s)$	Number of full-dimensional simplices met by s; number of internal facets crossed
$\mu,\ m,\ M$	Sampling law on K or V and its density bounds in (1.1)
$L_V,\ L_n$	Lower-hull heights of the lifted manifold V and of the n sites
$g_B,\ Z_B$	Power $\lvert p-c\rvert^2 - R^2$ of an empty ball B and contacts $\{p \in V : g_B(p) = 0\}$
$W,\ \mathcal{R},\ G$	Span of contact tangents and chords; regular set; open set in Theorem 1.1
$\mathcal{H}^{d-1}$	$(d-1)$-dimensional Hausdorff measure (zero on avoided exceptional sets)
$U_t(q)$	Thin axial query window around q (Section 4, Appendix B)
$H_u,\ P_\tau$	Screw motion and moving-sphere power (Section 5, Appendix C)

In either setting, the Delaunay triangulation lives in $\mathbb{R}^d$, so its full-dimensional simplices have $d+1$ vertices and facets of dimension $d-1$. The variable $N_n(s)$ counts every full-dimensional Delaunay simplex whose interior meets s. Whenever a Delaunay cell has more than $d+1$ cospherical vertices, we subdivide it using only sampled vertices by a fixed measurable triangulation rule that agrees on shared lower-dimensional faces. Write P for the convex hull of the n sites; if $\dim P < d$, no full-dimensional cell exists and $N_n(s) = 0$. For the crossing identity, suppose s has positive length, neither endpoint lies on a facet, and every facet met by s is crossed transversely through its relative interior (its interior within the facet's own affine hyperplane). Because P is convex, $s \cap \operatorname{int} P$ is a single interval whenever it is nonempty. Thus the segment enters $\operatorname{int} P$ at most once and gains one new simplex at each internal facet crossing between two simplices. Writing $\Xi_n(s)$ for the number of these internal facet crossings gives

$$N_n(s) = \mathbf{1}\{s \cap \operatorname{int} P \neq \varnothing\} + \Xi_n(s), \tag{2.1}$$

where the indicator $\mathbf{1}\{s \cap \operatorname{int} P \neq \varnothing\}$ equals one if s enters $\operatorname{int} P$ and zero otherwise, accounting for the first simplex visited. Every upper bound below counts intersections with simplex interiors as defined here, and every lower bound verifies that s passes through the interiors of the counted simplices.

2.1 Sites sampled in the volume

When sites fill a compact convex body $K \subset \mathbb{R}^d$, their typical inter-site spacing is of order $n^{-1/d}$. How close to the boundary ∂K can a fixed segment s lie while still visiting at most $1 + C_A|s|n^{1/d}$ cells in expectation? Two simple geometric facts yield the bound. First, for any fixed facet shape, the translations that make the facet meet s fill a region whose volume is proportional to $|s|$ and to the $(d-1)$-dimensional projected size of the facet. Second, a simplex of diameter r can cross s only if its open circumball is empty of other sites. Even when that circumball bulges outside K, the boundary margin $\operatorname{dist}(s, \partial K)$ lets us place a wide empty cone inside both the circumball and K, making large simplices exponentially unlikely once r exceeds $n^{-1/d}$.

Proposition 2.1. *Fix an integer $d \geq 1$ and a compact convex body $K \subset \mathbb{R}^d$ with nonempty interior. For each $n \geq 1$, draw n sites independently from a common probability law on K satisfying the volume-normalized analogue of* (1.1)*: every measurable subset $E \subset K$ occupying a fraction v of $\mathrm{vol}_d(K)$ receives probability between $m_0 v$ and $M_0 v$, for fixed constants $0 < m_0 \leq 1 \leq M_0 < \infty$. For every $A > 0$, there is a constant $C_A > 0$ depending only on A, d, K, m_0, M_0 such that every segment $s \subset K$ fixed before sampling satisfies*

$$\mathbb{E}N_n(s) \leq 1 + C_A|s|\, n^{1/d} \qquad \text{whenever } \operatorname{dist}(s, \partial K) \geq An^{-1/d}.$$

In particular, for each compact $Q \subset \operatorname{int} K$, a single constant works for all $n \geq 1$ and all fixed segments $s \subset Q$.

Proof. If $|s| = 0$, the segment meets at most one simplex interior, and if $n \leq d$ there are no full-dimensional simplices at all. We may therefore assume $\ell = |s| > 0$ and $n \geq d + 1$. Set $m = m_0/\mathrm{vol}_d(K)$ and $M = M_0/\mathrm{vol}_d(K)$, so every measurable region $W \subset K$ of volume w has probability between mw and Mw. In dimension $d = 1$, each interval visited after the first requires a sampled site inside s, and the expected number of sites in s is at most $Mn\ell$.

Suppose now that $d \geq 2$. Because the sampling law is absolutely continuous with respect to Lebesgue measure, almost surely no $d + 1$ sites lie in a hyperplane and no $d + 2$ lie on a sphere, so the Delaunay triangulation is unique and simplicial. Moreover, since s is fixed in advance, almost surely neither endpoint lies on a hyperplane spanned by d sites and the line through s misses every $(d - 2)$-dimensional affine subspace spanned by $d - 1$ sites. Thus s crosses every facet it meets transversely through its relative interior. By (2.1), even if part of s leaves the sample hull, every simplex met after the first requires an internal facet crossing. Letting I count all pairs consisting of a Delaunay simplex and one of its facets meeting s, we obtain $N_n(s) \leq 1 + I$.

To bound $\mathbb{E}I$, parametrize a candidate $(d + 1)$-tuple of sites by a base vertex $x \in \mathbb{R}^d$ and displacement vectors $z_1, \ldots, z_d \in \mathbb{R}^d$, so the vertices are $x, x + z_1, \ldots, x + z_d$, and write $r = \max_i |z_i|$. Its facet opposite $x + z_d$ is $x + F$, where $F = \operatorname{conv}(0, z_1, \ldots, z_{d-1})$. Holding $z_1, \ldots, z_d$ fixed, this facet meets s precisely when x lies in the Minkowski difference $s - F = \{u - v : u \in s,\ v \in F\}$, shown in Fig. 2(a). Let e be a unit vector parallel to s, let κ_j be the volume of the unit ball in $\mathbb{R}^j$, and write $\mathrm{proj}_{e^\perp}$ for orthogonal projection onto the hyperplane perpendicular to e. Slicing $s - F$ by lines parallel to e gives

$$\mathrm{vol}_d(s - F) = \ell\ \mathrm{vol}_{d-1}(\mathrm{proj}_{e^\perp} F) \leq \ell\, \kappa_{d-1} r^{d-1},$$

because each line above $\mathrm{proj}_{e^\perp} F$ sweeps out an interval of length ℓ when the projection is full-dimensional in $e^\perp$, while both sides vanish otherwise.

Next fix $z_1, \ldots, z_d$ and a base point $x \in s - F$, and pick an intersection point $y \in (x + F) \cap s$. The closed circumball $B(c, R)$ of the candidate simplex contains its facet $x + F$ and hence contains y. Because two vertices lie at distance at least r (if $r = |z_i|$, already x and $x + z_i$ are r apart), the circumradius satisfies $R \geq r/2$. Put $b = \operatorname{dist}(s, \partial K)$ and $a = \min(r, b)/4$. Moving a fraction a/R of the way from y toward c produces a point z with $|z - y| \leq a$ and $|z - c| \leq R - a$. Thus the open ball $B(z, a)$ lies inside the open circumball $\operatorname{int} B(c, R)$. Because every point of $B(z, a)$ stays within distance $2a \leq b$ of $y \in s$, convexity of K also places $B(z, a)$ inside K.

To turn this ball into a cone stretching across the simplex diameter, note that some pair of vertices is at least r apart. By the triangle inequality, a vertex v farthest from z satisfies $|v - z| \geq r/2$. Take the $(d - 1)$-dimensional disk of radius a through z perpendicular to $v - z$, and join it to v. By convexity, the interior of this cone of height $|v - z|$ (the shaded triangle in Fig. 2(b)) lies inside both $\operatorname{int} B(c, R)$ and K. Comparing its volume with the density lower bound m shows that a single site lands in the open circumball with probability p satisfying

$$p \ \geq\ \frac{m\kappa_{d-1}}{d}\, a^{d-1}|v - z| \ \geq\ c_0\, r \min(r, b)^{d-1}, \qquad c_0 = \frac{m\kappa_{d-1}}{2d\, 4^{d-1}}.$$

For the candidate simplex to be Delaunay, all $n-d-1$ remaining sites must miss this open circumball. Conditional on the candidate vertices, that emptiness event has probability $(1-p)^{n-d-1} \le e^{d+1}e^{-np}$.

Finally, there are $(d+1)\binom{n}{d+1} \le n^{d+1}/d!$ ways to choose an unordered $(d+1)$-tuple of candidate sites and distinguish one of its facets. Because the change of variables $(x, z_1, \dots, z_d) \mapsto (x, x+z_1, \dots, x+z_d)$ has Jacobian one, their joint density is at most M^{d+1}. Integrating over $x \in s - F$ contributes the factor $\ell\,\kappa_{d-1}r^{d-1}$. Meanwhile, as r varies, displacement tuples with $\max_i |z_i| \le r$ fill a product of d balls of radius r having volume $\kappa_d^d r^{d^2}$ and radial volume element $d^2\kappa_d^d r^{d^2-1}\,dr$. Multiplying by the spatial factor ℓr^{d-1} and integrating up to $\operatorname{diam} K$ yields

$$\mathbb{E}I \le C\ell n^{d+1} \int_0^{\operatorname{diam} K} r^{d^2+d-2} \exp\bigl[-c_0 nr \min(r,b)^{d-1}\bigr]\,dr.$$

Rescaling lengths via $t = n^{1/d}r$ and using the margin $bn^{1/d} \ge A$ bounds this by

$$C\ell n^{1/d} \int_0^{\infty} t^{d^2+d-2} \exp\bigl[-c_0 t \min(t,A)^{d-1}\bigr]\,dt.$$

For each fixed $A > 0$, the integral converges because for $t \ge A$ the exponential decays as $e^{-c_0 A^{d-1}t}$. Combining this bound on $\mathbb{E}I$ with $N_n(s) \le 1 + I$ proves the inequality; for a compact set $Q \subset \operatorname{int} K$, simply set $A = \operatorname{dist}(Q, \partial K) > 0$. □

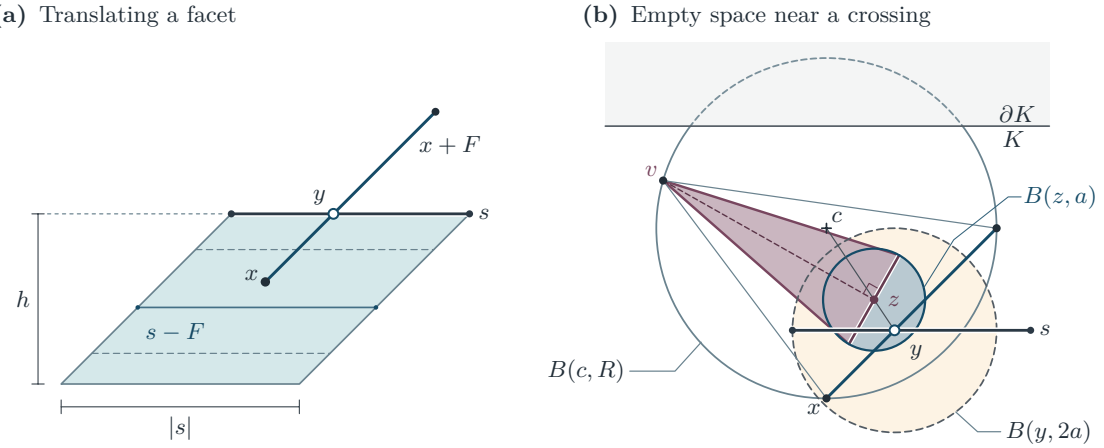


Figure 2: Counting crossings in a solid body. In (a), translating the facet F by x makes $x+F$ meet s precisely when x lies in the parallelogram $s-F$; every horizontal section through its interior has length $|s|$, and its height h is the length of the projection of F perpendicular to s. In (b), the upper dashed arc of the circumcircle extends outside K, while the disk $B(z,a)$ sits inside both the circumdisk and $B(y,2a) \subset K$. Joining a diameter of $B(z,a)$ to the distant vertex v gives the shaded triangle, the planar version of the empty cone used in the proof.

2.2 Reading the contacts from the lower hull

We now turn to sites sampled on a manifold $V \subset \mathbb{R}^d$, where an interior walk can lie far from every vertex of the cells it crosses. For instance, in Fig. 3(a), a horizontal segment inside an ellipse of 40 boundary sites starts inside one triangle and crosses 18 internal edges, visiting 19 triangles in all. To see how the surrounding curve dictates the shape of those triangles, we lift the geometry to $\mathbb{R}^{d+1}$, where empty spheres become supporting hyperplanes.

Lift each point $p \in V$ vertically onto the standard paraboloid $(p, |p|^2) \in \mathbb{R}^d \times \mathbb{R}$, where $|p|$ denotes the Euclidean norm of p. The *lower hull* of the lifted manifold is the lower boundary

of the convex hull $\operatorname{conv}\{(p, |p|^2) : p \in V\}$. Over each query $q \in K$, its height is the minimum vertical height attained inside that convex hull,

$$L_V(q) = \min\{h : (q, h) \in \operatorname{conv}\{(p, |p|^2) : p \in V\}\}. \tag{2.2}$$

Because V is compact, the lifted convex hull is compact as well, so the minimum is always attained. Repeating this lift for just the n sampled sites produces a piecewise affine convex function $L_n \geq L_V$ on their convex hull P. Each full-dimensional lower facet of the lifted sample hull projects orthogonally onto a Delaunay cell in $\mathbb{R}^d$, and subtracting any affine function from all lifted heights preserves those projected cells (see Brandt and Weinstein [8, Definitions 2.3 and 2.5 and Proposition 2.6 of the arXiv version] for the analogous lifting on a real algebraic variety).

At any interior point $q \in \operatorname{int} K$, choose an affine function $\ell\colon \mathbb{R}^d \to \mathbb{R}$ satisfying $\ell(q) = L_V(q)$ and $\ell(x) \leq L_V(x)$ throughout K. Its graph is a lower supporting hyperplane of the lifted hull at $(q, L_V(q))$. Write the linear part of ℓ as $2c \cdot p$ for a vector $c \in \mathbb{R}^d$. Completing the square turns the vertical gap $|p|^2 - \ell(p)$ into the power of a ball B with centre c and squared radius $R^2 = |c|^2 + \ell(0)$:

$$g_B(p) = |p|^2 - \ell(p) = |p - c|^2 - R^2 \geq 0, \qquad p \in V. \tag{2.3}$$

Since $|p - c|^2 - R^2$ is negative inside B, zero on ∂B and positive outside, $g_B(p) \geq 0$ means precisely that the open ball $\operatorname{int} B$ misses V. The supporting hyperplane meets the lifted manifold exactly at the lifts of the *contacts* $Z_B = \{p \in V : g_B(p) = 0\}$. Moreover, if $(q, L_V(q))$ is written as a convex combination of lifted points $(p_i, |p_i|^2)$ on V, their gaps $g_B(p_i) \geq 0$ must average to zero. Every point p_i with positive weight therefore satisfies $g_B(p_i) = 0$, so projecting back to $\mathbb{R}^d$ shows that $q \in \operatorname{conv} Z_B$.

This correspondence is completely reversible. Call a closed ball B *supporting at* q if its open interior misses V and the convex hull of its contacts Z_B contains q. Every lower supporting hyperplane at $(q, L_V(q))$ defines such a ball via (2.3). Conversely, given a supporting ball B at q, its affine function $\ell(p) = |p|^2 - g_B(p)$ lies below the paraboloid on V and meets it on Z_B. Expressing q as a convex combination from Z_B and lifting with those weights shows that $\ell(q) = L_V(q)$. Notice that a supporting ball may extend outside K. Whenever $q \notin V$, at least two distinct contacts are needed to contain q in their convex hull, whereas at $q \in V$ the radius-zero ball centred at q is itself a supporting ball.

Figure 3 visualizes this dictionary on the 40-site ellipse walk. Letting ℓ_q denote the supporting line at the midpoint q of s, the segment inside the shaded triangle of Fig. 3(b) corresponds to the highlighted affine segment of L_n in Fig. 3(c). Figure 3(d) then magnifies the vertical excess $L_n - L_V$, turning each triangle visited along s into a single smooth arch between crossings.

3 From contact geometry to a bound on the walk

Why should the spacing of sampled sites on a manifold V govern a Delaunay walk that travels far from V? Take a ball of small radius t in a coordinate chart of V. Its coordinate volume is proportional to t^δ, and on a small enough chart the Jacobian converting coordinate volume to intrinsic volume lies between two positive constants. Covering compact V by finitely many such charts shows that each resulting patch occupies a fraction of total intrinsic volume of order t^δ. Under (1.1), the probability that one site lands in the patch, called its *sampling mass*, is likewise of order t^δ. At scale $t = n^{-1/\delta}$, a patch therefore holds order one site on average. Subsection 3.1 formalizes a witness-and-collector lemma showing that if sites in a few witness patches force all vertices of nearby simplices into a small collector region, the expected local cell count stays bounded. Subsection 3.2 identifies the contact conditions on V that

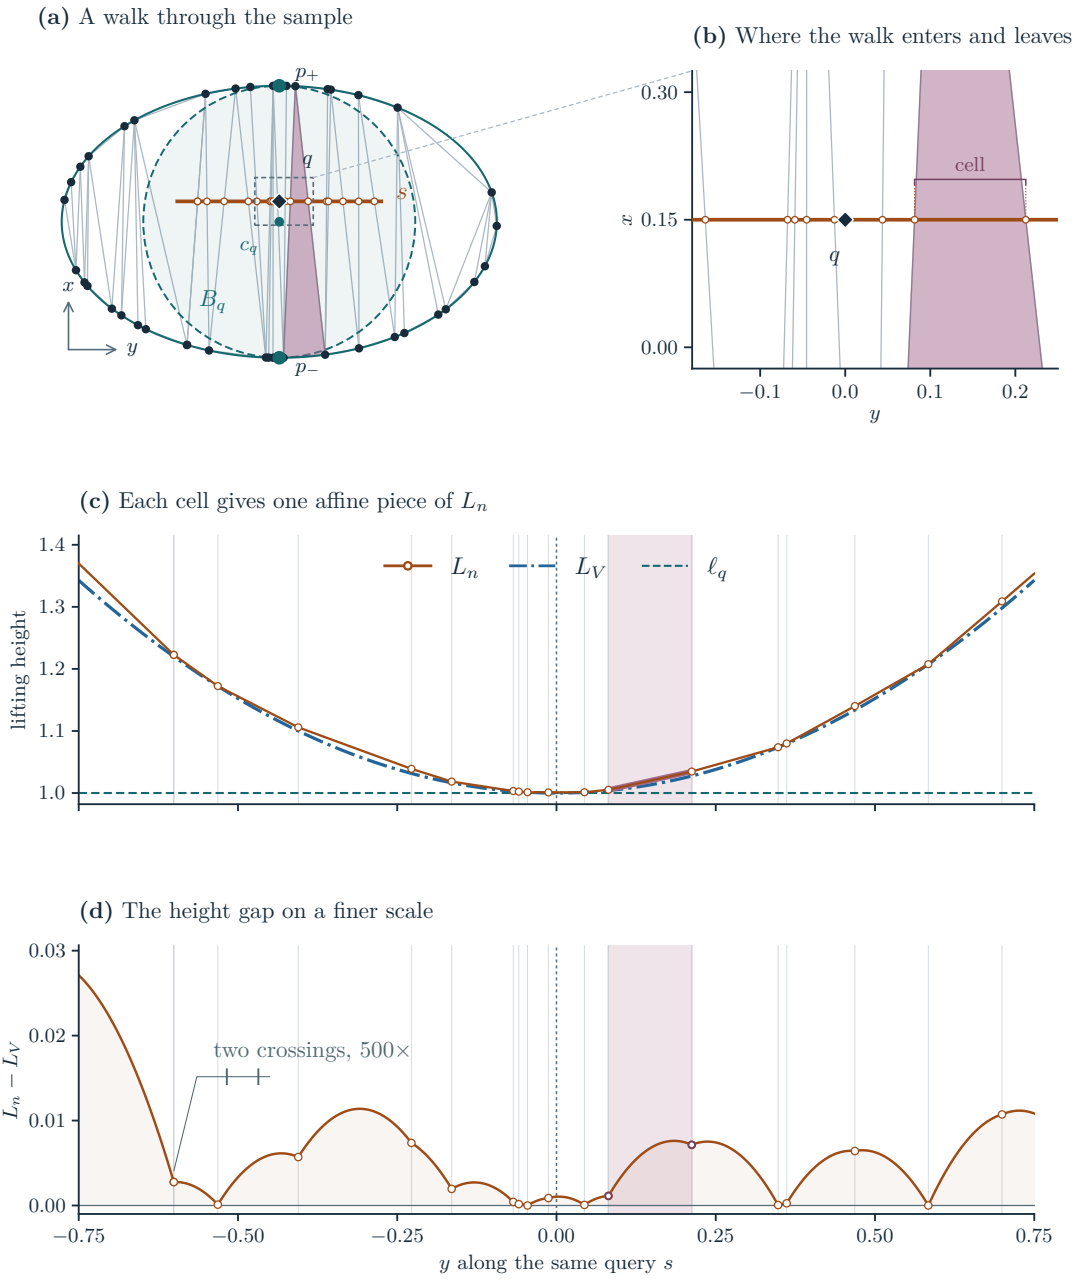


Figure 3: One walk in the Delaunay triangulation of 40 sites sampled by arclength on $x^2 + y^2/1.6^2 = 1$. In (a), y is horizontal, the segment runs from $y = -0.75$ to 0.75 at $x = 0.15$, and its midpoint q lies on the chord between the two points p_-, p_+ where the supporting circle (centred at c_q) touches the ellipse. In (b), the open circles mark where the segment enters and leaves the shaded triangle. In (c), follow the corresponding affine piece of the sample height L_n above $L_V = 1 + (1 - 1/1.6^2)y^2$, where ℓ_q is tangent to L_V at q. In (d), each arch of the excess $L_n - L_V$ corresponds to one triangle visited; the inset magnifies only the horizontal spacing of two of the 18 crossings near $y = -0.60$ by a factor of 500.

make this localization work both at regular queries and across codimension-one switches, and Subsection 3.3 illustrates both regimes on explicit embeddings.

3.1 Counting with occupied patches

Because no Delaunay simplex can hold a sampled site inside its open circumball, placing sites in a few well-chosen patches near a query q constrains where any nearby simplex's vertices can sit. In the witness-and-collector framework [14, Section 1.2], we fix a bounded number of deterministic *witness* patches on V and a small *collector* region. Whenever every witness patch receives at least one site, every vertex of every simplex meeting a small ball around q must fall inside the collector. At the single scale $t = n^{-1/\delta}$, a witness patch still stays empty with constant probability. To control the expectation, we test a continuous family of larger scales $v = u/n$. As u grows, the collector mass scales polynomially in u, whereas the chance of missing a witness patch decays exponentially in u. This same lemma will later power the counterexample bounds in Section 4.

Before stating the lemma formally, consider again the ellipse walk of Fig. 3. In Fig. 4(a), picking one site from each of four witness arcs flanking p_+ and p_- forms a quadrilateral enclosing $B(q, 2r)$. Any triangle reaching the smaller disk $B(q, r)$ cannot swallow all four witness sites inside its circumcircle. As Fig. 4(b) shows, all seven triangles meeting $B(q, r)$ therefore draw their vertices from two short collector arcs near p_+ and p_-; the rightmost purple interval in the enlarged horizontal diameter strip corresponds to the highlighted triangle of Fig. 3(b). Counting triples of sites on those two arcs bounds the local triangle count, and because this vertex-confinement argument depends only on the cell vertices, it survives the affine coordinate changes of Section 4.

Lemma 3.1. *For integers $d, n \geq 1$, draw n sites independently from a common probability measure μ on $\mathbb{R}^d$. In any measurable triangulation with vertices among these sites, let M_n count the closed d-simplices that meet a fixed set $D_n \subset \mathbb{R}^d$, and assume this count is measurable. Fix $a, A, v_0 > 0$ and a positive integer J. Suppose that at every scale $1/n \leq v \leq v_0$, there exist at most J deterministic measurable patches, each of μ-mass at least av, and a deterministic measurable region of μ-mass at most Av. Assume that whenever all these patches contain sampled sites, every vertex of every counted simplex lies in the region. Then $\mathbb{E}M_n \leq C(d, a, A, J, v_0)$.*

Proof. Fix $1 \leq u \leq nv_0$ and set $v = u/n$. A patch of μ-mass at least av is missed by all n independent sites with probability at most $(1 - av)^n \leq e^{-nav} = e^{-au}$. Meanwhile, the number Y of sites falling in the collector region is binomial with mean at most $n(Av) = Au$, so $\mathbb{E}e^Y \leq \exp((e-1)Au)$, and Markov's inequality gives $\mathbb{P}\{Y > ((e-1)A+1)u\} \leq e^{-u}$. Whenever every witness patch is occupied and $Y \leq ((e-1)A+1)u$, at most $\binom{Y}{d+1} \leq Y^{d+1}$ simplices can be formed from sites in the region. Setting $B_0 = ((e-1)A+1)^{d+1}$ and taking a union bound over the at most J empty-patch events and the upper tail of Y yields

$$\mathbb{P}\{M_n > B_0 u^{d+1}\} \leq Je^{-au} + e^{-u} \qquad (1 \leq u \leq nv_0).$$

Now put $B = \max\{B_0, v_0^{-(d+1)}\}$. Whenever $u > nv_0$, we have $Bu^{d+1} > n^{d+1} \geq M_n$, so $\mathbb{P}\{M_n > Bu^{d+1}\} = 0$ and the tail bound holds for all $u \geq 1$:

$$\mathbb{P}\{M_n > Bu^{d+1}\} \leq Je^{-au} + e^{-u} \qquad (u \geq 1).$$

Integrating this stretched-exponential tail over u bounds $\mathbb{E}M_n$ by a constant depending only on d, a, A, J, v_0. The same tail integration also bounds every higher moment $\|M_n\|_p = (\mathbb{E}|M_n|^p)^{1/p} \leq C(d, a, A, J, v_0, p)$ for $1 \leq p < \infty$. □

In particular, for the sampling law and open set G of Theorem 1.1, fix $1 \leq p < \infty$ and a compact subset $Q \subset G$. Applying Minkowski's triangle inequality for $\|\cdot\|_p$ to the sum of cell counts over the covering balls gives $\|N_n(s)\|_p \leq C_{Q,p}(1 + |s|n^{1/\delta})$ for every segment $s \subset Q$ fixed before sampling and all $n \geq n_Q$.

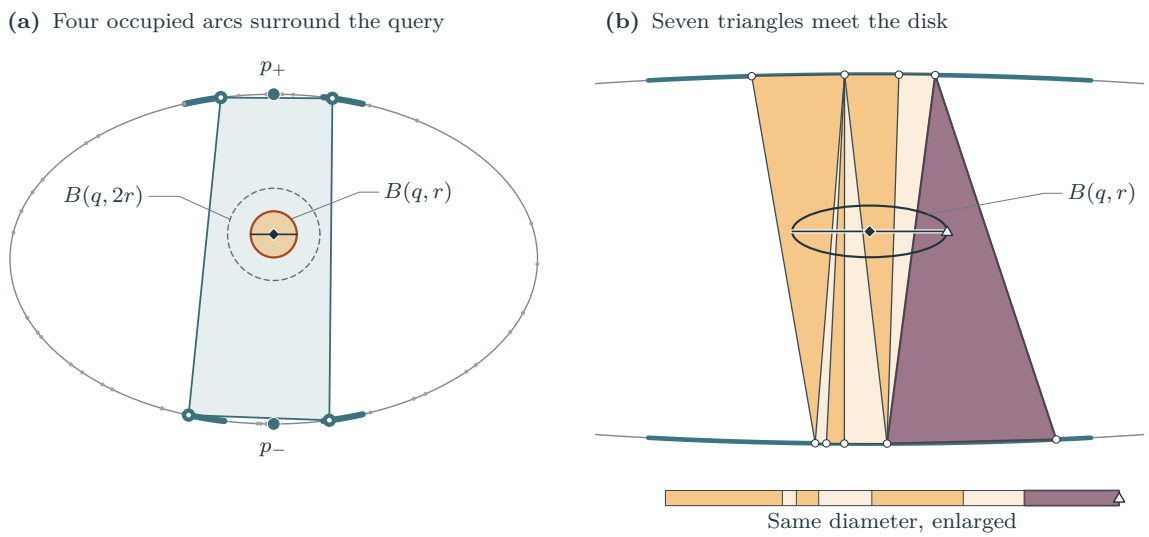


Figure 4: The sample of Fig. 3, now viewed near $q = (0.15, 0)$ with $r = 0.14$. In (a), picking one site from each of four witness arcs around the two contacts forms a quadrilateral enclosing $B(q, 2r)$; as in Fig. 3(a), a dark diamond marks q and filled circles mark the contacts. In (b), all seven triangles meeting $B(q, r)$ draw their vertices from four sites on the upper arc and five on the lower arc (all with $|y| \le 0.4$). The strip below the drawing enlarges the horizontal diameter of $B(q, r)$ across those seven triangles, with an open triangle marking its right endpoint at both scales and the rightmost purple interval belonging to the triangle highlighted in Fig. 3. Both views use a horizontal y-axis, stretched threefold in (b).

3.2 Which directions the contacts cover

To turn supporting balls into witness patches and collectors, the power of a supporting ball must grow cleanly away from each contact, and the directions available at those contacts must span $\mathbb{R}^d$. Call a contact point p of a supporting ball with centre c *ordinary* if the squared-distance function $x \mapsto |x - c|^2$, written in local coordinates on V, has a positive definite Hessian at p, meaning that V curves strictly away from the sphere in every tangent direction. Because every contact minimizes this function on V, its Hessian is always positive semidefinite; we call the contact *degenerate* if the Hessian has a zero eigenvalue. Near an ordinary contact, moving distance t along V increases the power by order t^2, which traps distant cell vertices.

Next examine the directions around q generated by a supporting ball B at q. Choose a set of contacts of B whose convex hull contains q, and call this representation *minimal* if no proper subset still contains q in its convex hull. Its contact points $p_1, \dots, p_j$ are affinely independent, so $q = \sum_{i=1}^{j} \lambda_i p_i$ with unique strictly positive weights $\lambda_i > 0$ summing to one. Moving any contact p_i slightly along V shifts this convex combination along a vector in the tangent space $T_{p_i}V$, while keeping contacts fixed and shifting weights moves it parallel to the chords $p_i - p_1$. Together, these infinitesimal moves span the linear subspace

$$W = \operatorname{span}\big(T_{p_1}V, \dots, T_{p_j}V,\, p_2 - p_1, \dots, p_j - p_1\big),$$

where each tangent space is translated to the origin. Call an interior query $q \in \operatorname{int} K$ *regular* if $q \notin V$, every contact of every supporting ball at q is ordinary (even contacts unused in a particular representation), and every minimal representation by contacts of any supporting ball at q satisfies $W = \mathbb{R}^d$. Write $\mathcal{R}$ for the set of all regular queries.

Theorem 3.2. *Fix integers $d \ge 2$ and $1 \le \delta \le d - 1$. Let $V \subset \mathbb{R}^d$ be a compact C^2 embedded δ-manifold without boundary and of full affine span. Write $K = \operatorname{conv} V$ and let $\mathcal{R}$ be the set of*

regular queries for V. Draw n sites independently on V from a common law satisfying (1.1). *For every compact $Q \subset \mathcal{R}$ there are constants $C_Q > 0$ and $n_Q \in \mathbb{N}$, depending only on V, Q and the bounds in* (1.1), *such that every segment $s \subset Q$ fixed before sampling satisfies*

$$\mathbb{E}N_n(s) \le C_Q(1 + |s|n^{1/\delta}) \qquad (n \ge n_Q).$$

If $(\operatorname{int} K) \setminus \mathcal{R}$ has zero $(d-1)$-dimensional Hausdorff measure, this estimate holds for $s = [a, b]$ for Lebesgue-almost every ordered pair $(a, b) \in (\operatorname{int} K)^2$, with constants that may depend on the pair.

Appendix A.1 proves Theorem 3.2. There, the spanning condition $W = \mathbb{R}^d$ lets sites at distance $O(t)$ from the contacts surround a ball of radius proportional to t around q, and quadratic growth of power confines every cell vertex into $O(t)$-balls around the local minima of power. Moreover, if one fixed ball supports every point along $s \subset Q \subset \mathcal{R}$, the same witness and collector patches serve the whole segment simultaneously. Appendix A.1 shows that $\mathbb{E}N_n(s) \le C_Q$ for all $n \ge n_Q$ uniformly over such segments.

Comparing the sample height with a support. Both Appendix A.1 and our later analysis of contact switches and axial slices rest on a simple height-comparison lemma for lower convex hulls. Assign each site x in a finite set a real height $H(x)$, and write L_n for the piecewise affine height of the lower hull of the lifted sites $(x, H(x))$. Over any full-dimensional projected cell, L_n agrees with an affine function ϕ. Call an affine function $\ell\colon \mathbb{R}^d \to \mathbb{R}$ a *support* if $\ell(x) \le H(x)$ at every site x, and define its *gap* at x by $g(x) = H(x) - \ell(x) \ge 0$. In our Delaunay lift $H(x) = |x|^2$, a supporting ball B defines an affine support ℓ_B via (2.3) whose gap $g_B(x)$ is precisely the power of B at x.

Suppose L_n stays within vertical height M of a support ℓ across a ball $B(x_0, r)$. Because $\phi \ge \ell$ at x_0 yet $\phi \le \ell + M$ on the entire sphere of radius r around x_0, the affine cell height ϕ can tilt relative to ℓ by slope at most M/r, as shown in Fig. 5(a). Near a point where the supporting ball switches from one contact set to another, L_n instead stays within height M of the maximum $\max(\ell_-, \ell_+)$ of two extreme supports. Lemma 3.3(c) shows that ϕ still stays within slope M/r of a suitable convex combination of ℓ_- and ℓ_+ (Fig. 5(b)), which will later confine cell vertices even when $B(x_0, r)$ straddles an interface of contact switch.

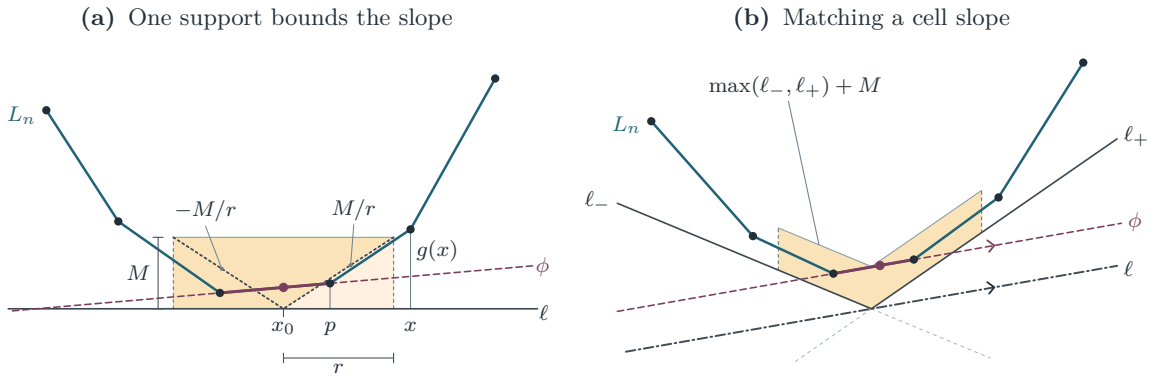


Figure 5: How a height bound controls the slope of a cell. In (a), subtracting the support ℓ places L_n in a strip of height M over an interval of radius r, limiting the cell slope to $\pm M/r$. In (b), the upper bound follows the higher of ℓ_- and ℓ_+; arrows mark the cell line ϕ and a convex combination ℓ of the supports, whose slopes match in this illustration while the lemma allows a discrepancy of at most M/r.

Lemma 3.3. *Fix an integer $d \ge 1$ and a finite set of sites in $\mathbb{R}^d$ with full affine span and assigned real heights H. Let x_0 lie in a full-dimensional cell of the projected lower hull, and*

let ϕ be the affine height on that cell. Fix $M \geq 0$. In (b) and (c), assume that the closed ball $B(x_0, r)$ of radius $r > 0$ lies in the convex hull of the sites.

(a) Let ℓ be a support and p a vertex of the cell. If $L_n \leq \ell + M$ at the point $x_0 + \lambda(p - x_0)$ for some $0 < \lambda \leq 1$, then $g(p) \leq M/\lambda$.

(b) If ℓ is a support and $L_n \leq \ell + M$ on $B(x_0, r)$, then $|\nabla(\phi - \ell)| \leq M/r$. In particular, $g(p) \leq g(x) + (M/r)|p - x|$ for every cell vertex p and every sampled site x.

(c) If $L_n \leq \max(\ell_-, \ell_+) + M$ on $B(x_0, r)$ for two supports ℓ_-, ℓ_+, there is a convex combination ℓ of them for which both conclusions of (b) hold.

Proof. For any support ℓ, set $f = \phi - \ell$. Because ϕ defines a lower supporting face of the lifted sites, we have $\phi \leq L_n$ on the sample hull and $\phi(x_0) = L_n(x_0) \geq \ell(x_0)$, whence $f(x_0) \geq 0$. At any site x, we have $\phi(x) \leq H(x)$ with equality at cell vertices, so $f(x) \leq g(x)$ with equality at vertices of the cell. In (a), affinity of f therefore gives $M \geq f(x_0 + \lambda(p - x_0)) = (1 - \lambda)f(x_0) + \lambda g(p)$; dropping the nonnegative first term yields $g(p) \leq M/\lambda$.

For (b) and (c), write $\Psi = \max(\ell_-, \ell_+)$, taking $\ell_- = \ell_+ = \ell$ in (b). We know $\phi(x_0) \geq \Psi(x_0)$ and $\phi(x_0 + re) \leq \Psi(x_0 + re) + M$ for every unit vector $e \in \mathbb{R}^d$. Pick an index $\sigma \in \{-, +\}$ attaining $\Psi(x_0 + re) = \ell_\sigma(x_0 + re)$. Using $\Psi(x_0) \geq \ell_\sigma(x_0)$ gives $\Psi(x_0 + re) - \Psi(x_0) \leq r\,\nabla\ell_\sigma \cdot e$, so comparing the inequalities at x_0 and $x_0 + re$ and dividing by r yields

$$\nabla\phi \cdot e \leq \max\{\nabla\ell_- \cdot e, \nabla\ell_+ \cdot e\} + M/r \qquad \text{for every unit vector } e.$$

Now let γ be the orthogonal projection of $\nabla\phi$ onto the line segment between $\nabla\ell_-$ and $\nabla\ell_+$. If $\gamma \neq \nabla\phi$, choose the unit vector $e = (\nabla\phi - \gamma)/|\nabla\phi - \gamma|$. By convexity, neither endpoint $\nabla\ell_\pm$ has a larger inner product with e than γ does, so the inequality above forces $|\nabla\phi - \gamma| \leq M/r$. The same bound holds trivially when $\gamma = \nabla\phi$. Writing γ as a convex combination of $\nabla\ell_-$ and $\nabla\ell_+$ defines the convex-combination support ℓ with $\nabla\ell = \gamma$, so $|\nabla f| \leq M/r$. Finally, for any cell vertex p and any site x,

$$g(p) = f(p) \leq f(x) + \frac{M}{r}|p - x| \leq g(x) + \frac{M}{r}|p - x|. \qquad \square$$

Where can a segment $[a, b]$ leave the regular set $\mathcal{R}$ for a generic embedding? Any set of dimension at most $d - 2$ has zero $(d - 1)$-dimensional Hausdorff measure and is missed by $[a, b]$ for almost every pair $(a, b) \in (\operatorname{int} K)^2$ (page 35). This excludes all queries supported by a sphere with a degenerate contact, whose image in query space has dimension at most $d - 2$. Call a minimal representation *deficient* if its contact directions satisfy $W \neq \mathbb{R}^d$. Why does the divisibility condition $d = j(\delta + 1)$ single out the only deficient configurations a generic segment can hit? Choosing $j \geq 2$ contacts on a δ-manifold together with $j - 1$ convex weights provides $j\delta + (j - 1) = j(\delta + 1) - 1$ degrees of freedom, while requiring a common tangent hyperplane and sphere imposes a second bound $2d - j(\delta + 1) - 1$ (Subsection A.2). The smaller of the two equals $d - 1 - |d - j(\delta + 1)|$, which stays at most $d - 2$ unless $d = j(\delta + 1)$. A single contact ($j = 1$) represents only itself, so its queries lie in $V \cap \operatorname{int} K$, of dimension at most δ, which again reaches $d - 1$ only when $d = 1 \cdot (\delta + 1)$. Whenever $\delta + 1$ does not divide d (for instance, a curve in $\mathbb{R}^3$ or a surface in $\mathbb{R}^4$ or $\mathbb{R}^5$), Theorem 3.2 therefore already establishes Theorem 1.1 without any crossing analysis.

In the remaining equality case $d = j(\delta + 1)$, a generic segment can cross the resulting $(d - 1)$-dimensional hypersurface, which we call the *interface*. At a typical point q of the interface, the j contacts representing q, called the *base contacts*, provide a hyperplane W of dimension $d - 1$. Equidistance from the j base contacts and tangency along their δ-dimensional tangent spaces impose $j - 1 + j\delta = d - 1$ independent linear constraints on the sphere centre. All tangent sphere centres therefore fill a one-dimensional line perpendicular to W, and emptiness

cuts out a closed interval on this line. For $j = 1$ on a nonconvex surface in $\mathbb{R}^3$, a segment inside $\operatorname{int} K$ simply crosses V at q. There the centres sweep the normal line from the largest empty ball on one side of T_qV, through the radius-zero ball at q, to the largest empty ball on the opposite side. For $j = 2$ on a curve in $\mathbb{R}^4$, Fig. 6 depicts this closed interval of centres explicitly. In both cases, the two endpoint spheres each gain one extra contact, on opposite sides of the affine hyperplane $q + W$. Sites sampled near those extra contacts supply the missing normal direction on their respective sides, allowing Lemma 3.3(c) to carry the local bound across the interface (Subsections A.3–A.5).

3.3 What changes when one direction is missing?

To see concretely how an interval of supporting centres opens up when one ambient direction is missing, we compare two explicit embeddings of the unit δ-sphere $S^\delta = \{(x, u) \in \mathbb{R} \times \mathbb{R}^\delta : x^2 + |u|^2 = 1\}$. Both embeddings keep the same two marked contacts and differ only by an extra coordinate. For any integer $\delta \geq 1$, we draw n independent sites on each embedded image from a law satisfying (1.1).

Two contacts that enclose an open set of queries. Consider first the map into $\mathbb{R}^{2\delta+1}$ given by

$$F(x, u) = (x, 2u, xu).$$

Because the first coordinate x and the middle block $2u$ determine (x, u), F is a smooth embedding. Its poles are $(x, u) = (\pm 1, 0)$, whose images are $(\pm 1, 0, 0)$. The identity $|F(x, u)|^2 - 1 = (3+x^2)|u|^2$ shows that the unit ball is empty of $F(S^\delta)$ and touches it only at those two pole images, with positive definite Hessian at both contacts. At $(\pm 1, 0, 0)$, tangent vectors have the form $(0, 2v, \pm v)$ for $v \in \mathbb{R}^\delta$. Thus the two tangent spaces span the last 2δ coordinate directions, while the chord between the pole images supplies the first direction. Hence $W = \mathbb{R}^{2\delta+1}$, making the origin 0 (the midpoint of the pole chord) a regular query.

Let us check that regularity persists on an open ball around 0. For points $p_1, p_2 \in F(S^\delta)$ near the two poles, demanding tangency at p_1 and p_2 imposes one equidistance equation $(p_1 - p_2) \cdot c = \frac{1}{2}(|p_1|^2 - |p_2|^2)$ and 2δ orthogonality equations $(p_i - c) \cdot T_{p_i} F(S^\delta) = 0$. At the poles, the coefficient vectors of these $2\delta + 1$ linear equations form a basis of $\mathbb{R}^{2\delta+1}$. A unique centre $c(p_1, p_2)$ therefore exists nearby and depends smoothly on (p_1, p_2), while positive Hessians and the positive gap off the poles keep the ball empty. By the same basis property, the map $(p_1, p_2, \lambda) \mapsto (1 - \lambda)p_1 + \lambda p_2$ has an invertible derivative at the poles with $\lambda = 1/2$. Every query q in a small neighbourhood of the origin therefore admits such a two-point representation with positive weights. Any supporting sphere at q must touch both points and agree with this unique sphere. Consequently, Theorem 3.2 bounds $\mathbb{E}N_n(s)$ by $O(1 + |s|n^{1/\delta})$ for every fixed segment s inside a small closed ball around the origin, giving $O(1 + n|s|)$ for a curve in $\mathbb{R}^3$ and $O(1 + \sqrt{n}\,|s|)$ for a surface in $\mathbb{R}^5$.

Two contacts that leave one direction free. Now keep the same poles and tangents in $\mathbb{R}^{2\delta+1}$, and append one coordinate using the first entry u_1 of u:

$$G(x, u) = (x, 2u, xu, u_1|u|^2) \in \mathbb{R}^{2\delta+2}.$$

Because the new coordinate vanishes to third order at the poles, the pole chord and both tangent spaces span only the hyperplane $\{x_{2\delta+2} = 0\}$ in $\mathbb{R}^{2\delta+2}$. Even so, the origin remains strictly inside $\operatorname{conv} G(S^\delta)$. Any supporting hyperplane through the origin would have to contain the pole chord and both tangent spaces and hence equal $\{x_{2\delta+2} = 0\}$, whereas $u_1|u|^2$ takes both signs on S^δ.

Since the origin is the midpoint of the two pole images, any sphere supporting the origin must touch both poles. Its centre therefore lies on the last coordinate axis, say at $te_{2\delta+2}$ with $t \in \mathbb{R}$. Passing through $(\pm 1, 0, 0, 0)$ forces squared radius $1 + t^2$. Substituting $|F(x,u)|^2 = 1 + (3 + x^2)|u|^2$ and $|u|^2 = 1 - x^2$ gives power

$$g_t(x,u) = |G(x,u) - te_{2\delta+2}|^2 - (1+t^2) = |u|^2\big[3 + x^2 + u_1^2(1-x^2) - 2tu_1\big].$$

At $t = 2$, the bracket factors as $(1-u_1)(3-u_1)+x^2(1-u_1^2) \geq 0$, which vanishes on $S^\delta \setminus \{u = 0\}$ only when $u_1 = 1$ (forcing $x = 0$ and $u = e_1$). By reflection $u_1 \mapsto -u_1$, the bracket at $t = -2$ vanishes only at $u = -e_1$. For intermediate $-2 < t < 2$, g_t is a strict convex combination of g_{-2} and g_2, so it vanishes exclusively at the two poles. For $|t| > 2$, the power turns negative at the extra contact $p_+ = (0, 2e_1, 0, 1)$ or $p_- = -p_+$. The supporting centres at the origin thus fill the closed interval $[-2, 2]$ along the last coordinate axis, as drawn for $\delta = 1$ in Fig. 6(a).

Figure 6(b) shows how the supporting sphere jumps across a segment through the origin. On the half-segment s_+ toward p_+, every point lies in the triangle spanned by the two poles and p_+, so the $t = 2$ sphere supports all of s_+; symmetrically, the $t = -2$ sphere supports s_-. Notice that at each pole $g_t = 4|u|^2 + O(|u|^3)$ uniformly over $t \in [-2, 2]$. Near p_+, writing $y = (x, u_2, \ldots, u_\delta)$ and $u_1 = \sqrt{1 - |y|^2}$ gives $g_2 = |y|^2 + O(|y|^4)$ (and symmetrically for g_{-2} near p_-). Both extra contacts are therefore ordinary, and $\partial_t g_t = -2u_1|u|^2$ is nonzero at $p_\pm$. Appendix A.4 produces a smooth hypersurface interface through the origin separating the two endpoint regimes. Proposition A.1 then gives the bound $O(1 + |s|n^{1/\delta})$ across a full open neighbourhood of the origin, including segments that cross the interface.

4 A circle of contacts and a linear walk

How drastically can a Delaunay walk change if we keep its length fixed and merely rotate the segment? On a prolate spheroid, an ellipsoid stretched along one axis of revolution, an empty ball can touch the surface along an entire circular latitude (Fig. 7(a)). A segment lying inside the planar disk bounded by one such contact circle uses only that single supporting ball, so it visits $O(1)$ cells in expectation. By contrast, a segment of the same length with nonzero displacement along the long axis crosses a continuum of contact disks and visits order n cells in expectation (Corollary 4.3). To explain this contrast and extend it to higher codimensions, we trace cell vertices back to thin axial belts along the manifold.

Let $E_{R,a}$ be the solid ellipsoid $x^2/R^2 + |y|^2/a^2 \leq 1$ in $\mathbb{R} \times \mathbb{R}^{d-1}$, where $0 < a < R$ and x is the coordinate along the long axis. Every boundary point $(x, y) \in \partial E_{R,a}$ has squared Euclidean norm $h(x) = a^2 + (1 - a^2/R^2)x^2$, depending on x alone. Subtracting the tangent line to h at any interior level $v \in (-R, R)$ leaves power $(1 - a^2/R^2)(x - v)^2$. Thus the supporting sphere touches $\partial E_{R,a}$ along the whole slice $\partial E_{R,a} \cap \{x = v\}$, called the *contact slice* at level v. Because the convex hull of this contact slice fills the entire cross-section $E_{R,a} \cap \{x = v\}$, one sphere supports every query point in that cross-section.

A contact slice can even have disconnected components while still supporting the full cross-section. Take the ellipse $\mathcal{E} = \{(x, y) : x^2/4 + y^2 = 1\}$ in $\mathbb{R}^2$ and form the surface $V = \mathcal{E} \times S^1 \subset \mathbb{R}^4$, with the unit circle S^1 in an orthogonal plane. For any site $p = (x, y, z)$ with $(x, y) \in \mathcal{E}$ and $z \in S^1$, the squared norm still depends only on x. Its height and power above the tangent line at level v equal

$$|p|^2 = h(x) = 2 + \tfrac{3}{4}x^2, \qquad h(x) - h(v) - h'(v)(x - v) = \tfrac{3}{4}(x - v)^2. \tag{4.1}$$

For $|v| < 2$, the contacts split into two disjoint circles at $y = \pm\sqrt{1 - v^2/4}$. Yet their convex hull still forms the solid cylinder filling the cross-section at $x = v$, drawn in Fig. 14(a) in Appendix B.

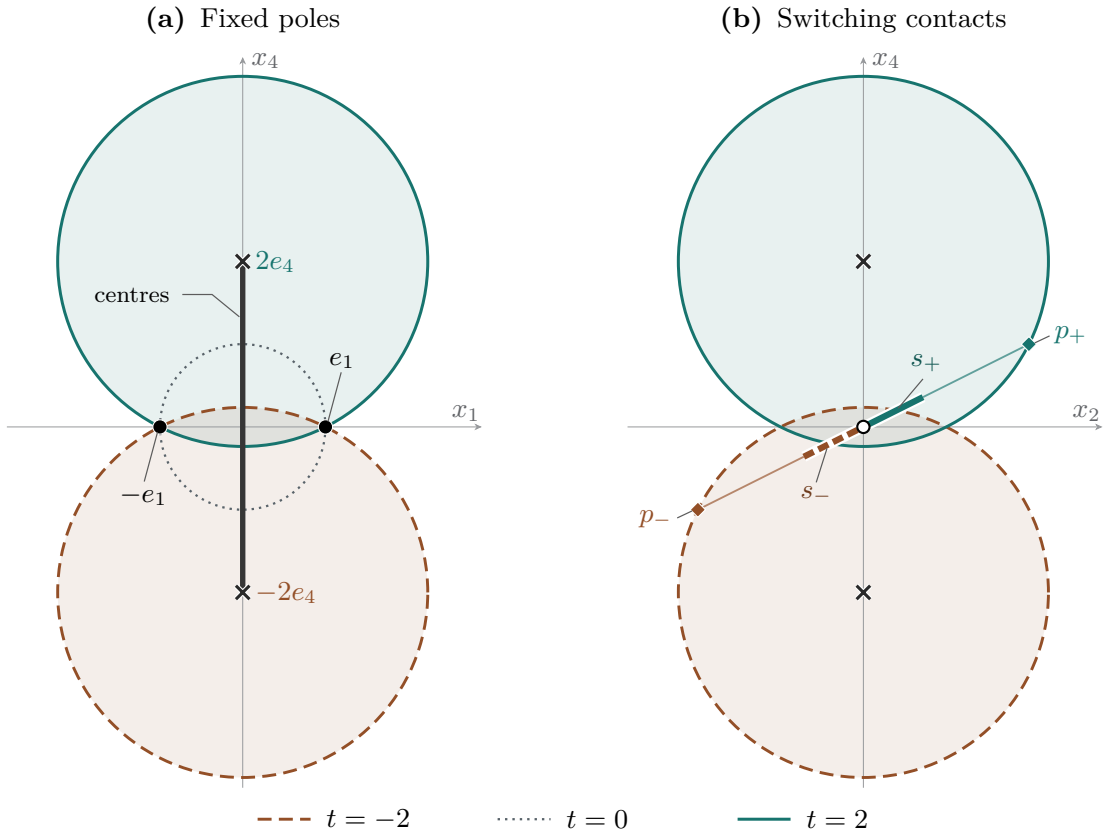


Figure 6: Two coordinate sections of the supporting-sphere family for G in $\mathbb{R}^4$ ($\delta = 1$), where e_i denotes the ith coordinate vector. In (a), the centres fill the vertical bar te_4 for $-2 \le t \le 2$, and each drawn sphere passes through both poles $\pm e_1$. In (b), the open circle marks the origin on the query segment: the dashed sphere centred at $-2e_4$ supports the dashed half s_- toward $p_- = (0, -2, 0, -1)$, while the solid sphere centred at $2e_4$ supports the solid half s_+ toward $p_+ = -p_-$. Notice that the poles in (a) lie outside the (x_2, x_4)-plane shown in (b).

In both examples, the lifted height $h(x(p))$ depends exclusively on the first coordinate $x(p)$. Allowing a general strictly convex axial height h lets us map ellipsoids of arbitrary aspect ratio a/R onto the unit sphere without losing track of Delaunay cells. Specifically, lift each sampled site p to $(p, h(x(p)))$ and project the lower faces of the lifted convex hull to $\mathbb{R}^d$. Refining any non-simplicial cells using only sampled vertices as in Section 2, let $N_n(s)$ count the resulting full-dimensional simplices whose interiors meet s. When $h(x(p)) = |p|^2$, this matches the ordinary Delaunay count.

Proposition 4.1. *Fix integers $d \ge 2$ and $1 \le \delta \le d - 1$. Let $V \subset \mathbb{R}^d$ be a compact C^2 embedded δ-manifold without boundary and of full affine span, and let x_-, x_+ be the minimum and maximum of its first coordinate x. Assume that*

1. *$h \in C^2([x_-, x_+])$ satisfies $0 < m \le h'' \le M$ on the entire interval, and*
2. *at every level $x_- < v < x_+$, every point of $V \cap \{x = v\}$ has a tangent vector to V with nonzero first coordinate, and $\operatorname{conv}(V \cap \{x = v\}) = (\operatorname{conv} V) \cap \{x = v\}$.*

Fix a compact set $Q \subset \operatorname{int} \operatorname{conv} V$ and an open set $I \subset (x_-, x_+)$ containing $x(Q)$. Draw n independent sites on V from a law giving probability zero to sets of intrinsic volume zero. On the strip $V \cap x^{-1}(I)$, assume that every measurable subset occupying a fraction r of total intrinsic volume receives probability between $\lambda_- r$ and $\lambda_+ r$ for fixed $0 < \lambda_- \le \lambda_+ < \infty$. Using

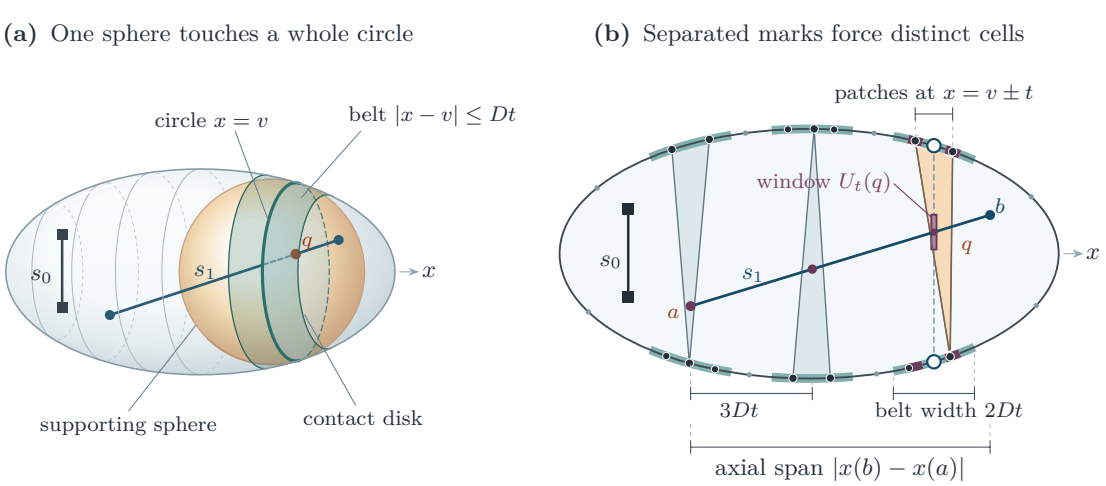


Figure 7: Axial slices and vertex localization. In (a), an empty sphere touches the prolate spheroid along a circle whose disk contains the query q: segment s_0 stays inside one axial slice, whereas s_1 crosses successive contact disks. In (b), a planar elliptic cross-section illustrates the localization lemma: around q at axial coordinate v, the rectangular window $U_t(q)$ has axial half-width ρt and transverse half-width ρ. Two short dark witness patches flank each contact (open circle) at $x = v$, forcing every Delaunay triangle meeting $U_t(q)$ to draw all its vertices from the belt $|x - v| \le Dt$. Three successful marks spaced $3Dt$ apart in axial coordinate along s_1 therefore have disjoint belts and lie in distinct triangles (scale t enlarged).

lifted height $h(x(p))$, suppose further that every affine hyperplane has sampling probability zero. Then every segment $s = [a, b] \subset Q$ fixed before sampling satisfies

$$c(1 + n|x(b) - x(a)|) \;\le\; \mathbb{E}N_n(s) \;\le\; C(1 + n|x(b) - x(a)|), \qquad n \ge n_0,$$

where $c, C > 0$ and $n_0 \in \mathbb{N}$ depend only on $V, Q, I, m, M, \lambda_-, \lambda_+$. The upper bound also holds without the hyperplane assumption.

Proof. Fix any query $q \in Q$ at axial level $x(q) = v$. For every small scale $0 < t \le t_0$, Lemma B.1 in Appendix B constructs a rectangular window $U_t(q)$ around q of axial radius ρt and transverse radius ρ, along with at most J witness patches on $V \cap x^{-1}(I)$ of sampling mass at least $c_1 t$. Whenever all witness patches are occupied, $U_t(q)$ lies inside the sample hull. Moreover, every vertex of every projected lower cell meeting $U_t(q)$ falls inside the belt $V \cap \{|x - v| \le Dt\}$, which has sampling mass $O(t)$. All constants c_1, ρ, D, t_0, J are uniform over $q \in Q$, and Fig. 7(b) sketches this geometry in 2D.

To prove the upper bound, take $n \ge \lceil 1/t_0 \rceil$ and fix the window $U_{1/n}(q)$. For every scale $1/n \le \tau \le t_0$, the larger window $U_\tau(q)$ contains $U_{1/n}(q)$. Thus the witness patches (mass $\ge c_1\tau$) and the collector belt (mass $O(\tau)$) at scale τ confine all cells meeting $U_{1/n}(q)$, and Lemma 3.1 bounds their expected count by a constant. Partitioning $s = [a, b]$ into equal pieces of axial length at most ρ/n and transverse length at most ρ (using boundedness of Q) covers s by $O_Q(1 + n|x(b) - x(a)|)$ windows centred at their midpoints, proving the upper bound.

For the lower bound, set $t = m_0/n$ with $m_0 > 0$ large enough that the probability of missing any witness patch at a mark is below $1/2$. All witness patches at a mark are then occupied with probability at least $1/2$. Starting at a, place marks along s at axial spacing $3Dt$ (using only a if $x(a) = x(b)$), and call a mark *successful* if all its witness patches contain sites. As verified in Appendix B, the hyperplane assumption ensures almost surely that each successful mark lies in the interior of a single simplex, whose vertices all lie in that mark's belt of axial half-width Dt. Because consecutive marks are separated by axial distance $3Dt$, their belts are pairwise disjoint, preventing any simplex from containing two successful marks. Summing probabilities

over the at least $1 + \lfloor |x(b) - x(a)|/(3Dt) \rfloor$ marks gives $\mathbb{E}N_n(s) \geq \frac{1}{2}(1 + \lfloor |x(b) - x(a)|/(3Dt) \rfloor)$, completing the proof. □

Notice that the hyperplane condition holds automatically whenever V is connected, real-analytic and of full affine span. Any affine hyperplane equation restricts to a non-identically-zero analytic function on V, whose zero set has intrinsic volume zero in local charts [26]. Replacing S^1 in $\mathcal{E} \times S^1$ by higher-dimensional spherical embeddings extends this linear-walk construction to every intrinsic dimension $2 \leq \delta \leq d - 1$.

Corollary 4.2. *For every pair of integers $d \geq 3$ and $2 \leq \delta \leq d - 1$, there is a compact connected real-analytic embedded manifold $V \subset \mathbb{R}^d$ of dimension δ, without boundary and of full affine span, to which Proposition 4.1 applies with the Delaunay lift. Draw n sites independently by normalized intrinsic volume on V. For each segment $s = [a, b] \subset \operatorname{int} \operatorname{conv} V$ fixed before sampling, $\mathbb{E}N_n(s)$ stays bounded as n grows if $x(a) = x(b)$, and is of order n if $x(a) \neq x(b)$. The latter case holds for Lebesgue-almost every ordered pair of interior endpoints.*

Appendix B (page 47) constructs the spherical factor and proves Corollary 4.2. Returning to the prolate spheroid, affine rescaling shows how linear growth depends on aspect ratio a/R uniformly as the spheroid stretches.

Corollary 4.3. *Fix an integer $d \geq 2$ and $0 < \kappa < 1$. For $0 < a < R$, draw n sites independently by normalized $(d-1)$-dimensional surface measure on the boundary of $E_{R,a} \subset \mathbb{R}^d$. Fix a segment $s \subset \kappa E_{R,a}$ before sampling, where the dilation is about the centre, and let Δx be the difference of the axial coordinates of its endpoints. Then*

$$c(1 + n|\Delta x|/R) \;\leq\; \mathbb{E}N_n(s) \;\leq\; C(1 + n|\Delta x|/R), \qquad n \geq n_0,$$

where $c, C > 0$ and $n_0 \in \mathbb{N}$ depend only on d, κ and hold uniformly for all $0 < a < R$.

Proof. The linear map $T = \operatorname{diag}(R, a, \ldots, a)$ carries the unit sphere S^{d-1} onto $\partial E_{R,a}$. For any point $\xi = (u, z) \in S^{d-1}$, the transformed site $T\xi$ has squared norm $|T\xi|^2 = a^2 + (R^2 - a^2)u^2$, so the vertical affine transformation

$$(p, H) \longmapsto (T^{-1}p,\, (H - a^2)/(R^2 - a^2))$$

maps lifted sites $(T\xi, |T\xi|^2)$ directly to (ξ, u^2). Because vertical upward rays map to vertical upward rays, lower faces map to lower faces, and T^{-1} transports both the Delaunay cells and our refinement rule without changing the crossing count. Figure 8 illustrates why using the transported height u^2 is essential: the ordinary paraboloid lift of S^{d-1} puts every site at height one and collapses the triangulation to a single cell.

We next compute how T^{-1} distorts surface measure, writing $\eta = a/R \in (0, 1)$ for the aspect ratio. Take an orthonormal tangent basis at $\xi \in S^{d-1}$ together with the outward unit normal ξ, and apply T. The image parallelepiped has volume $\det T = Ra^{d-1}$, and its height above the transformed tangent hyperplane equals $(T\xi) \cdot (T^{-1}\xi/|T^{-1}\xi|) = 1/|T^{-1}\xi|$. Thus T stretches surface area by $\det(T)|T^{-1}\xi| = Ra^{d-2} j_\eta(u)$, where $j_\eta(u) = \sqrt{1 - u^2 + \eta^2 u^2}$. Writing σ for normalized surface measure on S^{d-1} and cancelling Ra^{d-2} upon normalization, a single site falls in a measurable subset $A \subset S^{d-1}$ under T^{-1} with probability

$$\frac{\int_A j_\eta(u)\, d\sigma}{\int j_\eta\, d\sigma}, \qquad j_\eta(u) = \sqrt{1 - u^2 + \eta^2 u^2}.$$

The denominator stays bounded between the positive constant $\int \sqrt{1 - u^2}\, d\sigma$ and 1. Putting $\kappa' = (1 + \kappa)/2$, on the strip $|u| < \kappa'$ the integrand $j_\eta(u)$ lies between $\sqrt{1 - (\kappa')^2}$ and 1, uniformly in $\eta \in (0, 1)$.

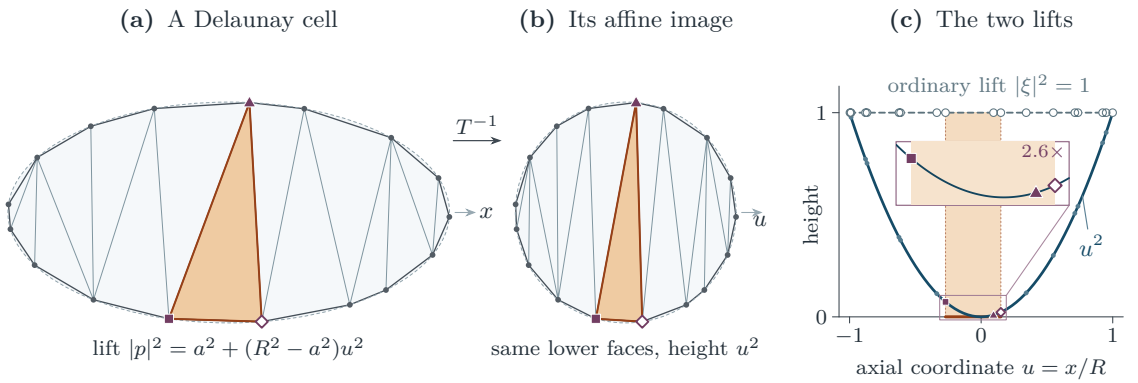


Figure 8: Transporting the lifted height from an ellipse to the unit circle. In (a), 17 sites on the ellipse with semiaxes 2 and 1 form a Delaunay triangulation; the triangle, square and open diamond mark the same three vertices across (a)–(c). In (b), the inverse scaling T^{-1} maps the ellipse to the unit circle while preserving the shaded triangle. In (c), plotting the transported height u^2 shows the nondegenerate lower hull of the same cells (enlarged 2.6× in the inset), whereas the ordinary circle lift $|\xi|^2 = 1$ collapses all sites onto height one.

Now apply Proposition 4.1 to the unit sphere with height $h(u) = u^2$, taking $Q = \kappa B^d$ (the closed ball of radius κ) and $I = (-\kappa', \kappa')$. Only the poles $u = \pm 1$ are critical for u. At every interior level $v \in (-1, 1)$, $\mathrm{conv}(S^{d-1} \cap \{u = v\})$ fills the entire cross-section of the unit ball (even for the two-point slice when $d = 2$). Moreover, every affine hyperplane has σ-measure zero, $h'' = 2$, and the density bounds on I depend only on d and κ. Because $T^{-1}s$ has axial displacement $|\Delta x|/R$, the proposition gives the claimed uniform bounds. (Notice that at $a = R$, all sites lie on one sphere and the height rescaling is undefined.) □

Finally, contrast the prolate spheroid with an ellipsoid having a *unique shortest* semiaxis (covering both triaxial ellipsoids as in Fig. 1 and oblate spheroids). Here the contact slices shrink back to pairs of ordinary points, restoring the intrinsic-dimension rate and collapsing to $O(1)$ along the shortest axis.

Proposition 4.4. *For any integer $d \geq 2$, let $K_* \subset \mathbb{R}^d$ be a solid ellipsoid with positive semiaxes $a_1, \ldots, a_d$ satisfying $a_d < a_i$ for $i < d$. Draw n sites independently by normalized $(d-1)$-dimensional surface measure on ∂K_*. Every interior point of K_* is a regular query. For each compact $Q \subset \operatorname{int} K_*$ there are $C_Q > 0$ and an integer n_Q such that every segment $s \subset Q$ fixed before sampling satisfies*

$$\mathbb{E}N_n(s) \leq C_Q(1 + |s|n^{1/(d-1)}), \qquad n \geq n_Q.$$

For segments parallel to the shortest axis, $\mathbb{E}N_n(s) \leq C_Q$ for $n \geq n_Q$. These constants may depend on the fixed semiaxes.

Proof. Write an interior point as $q = (w, z) \in \mathbb{R}^{d-1} \times \mathbb{R}$. The vertical line parallel to the shortest axis through q meets ∂K_* at $p_\pm = (w, \pm b(w))$, where $b(w) = a_d\sqrt{1 - \sum_{i<d} w_i^2/a_i^2} > |z|$. On ∂K_*, the relation $z'^2 = a_d^2(1 - \sum_{i<d} w_i'^2/a_i^2)$ makes the lifted height depend on w' alone:

$$|(w', z')|^2 = H(w') = a_d^2 + \sum_{i<d} c_i w_i'^2, \qquad c_i = 1 - a_d^2/a_i^2 > 0.$$

Because H is strictly convex in w', Jensen's inequality shows that any convex combination of lifted boundary points projecting to $q = (w, z)$ has height at least $H(w)$, with equality attained by interpolating between p_- and p_+ with the barycentric weights $(1 - z/b(w))/2$ and

$(1 + z/b(w))/2$ of q. Hence $L_{\partial K_*}(w, z) = H(w)$ throughout $\operatorname{int} K_*$. Since H is differentiable, $L_{\partial K_*}$ has a unique supporting hyperplane at q. Its gap at $(w', z') \in \partial K_*$ is

$$H(w') - H(w) - \nabla H(w) \cdot (w' - w) = \sum_{i<d} c_i (w'_i - w_i)^2,$$

vanishing exclusively at the two contacts p_+ and p_-.

At both contacts, the last coordinate $\pm b(w)$ is nonzero, so w' gives local coordinates on ∂K_* with positive definite Hessian $2\,\mathrm{diag}(c_1, \ldots, c_{d-1})$. Because the boundary normal at $p_\pm$ has nonzero vertical component, the chord $p_+ - p_-$ is transverse to $T_{p_\pm}\partial K_*$, and the tangent spaces and chord span $\mathbb{R}^d$. Moreover, q lies strictly between p_- and p_+, giving strictly positive weights in its unique minimal representation. Every $q \in \operatorname{int} K_*$ is therefore regular, and Theorem 3.2 yields the first bound. Along any vertical segment, the same two contacts and supporting ball serve every query point, so the common-ball estimate of Appendix A.1 lowers the bound to C_Q. □

5 A fourth-order contact on a space curve

For a curve in $\mathbb{R}^3$ ($d = 3$, $\delta = 1$), the generic spacing bound already allows linear growth $O(n)$. Can the expected number of crossed tetrahedra grow faster than linearly even when all contacts are isolated? Here we show that it can, starting from a reference family of moving empty spheres of constant radius 2 that touch three ideal helical arcs at only three isolated points: a front point B and two rear points R_+, R_-, with the query segment piercing their triangle BR_+R_- in Fig. 9(a). Measured in signed arclength along each arc, the sphere power grows quartically at B and quadratically at R_+ and R_- (Fig. 9(b)), allowing two sites with nearby phases, one near each rear contact to form a shared rear edge paired with a long run of consecutive sites near B. The smooth closed curve of Corollary 5.2 keeps these three helical arcs and constant-radius spheres exact over a fixed parameter interval, while the two-component real-analytic link of Corollary 5.3 realizes the same $(4, 2, 2)$ pattern of contact orders over a fixed interval using a different one-parameter family of spheres of varying radius; both feed into the same counting argument (Proposition D.2) to attain $\Omega(n^{4/3})$. By contrast, for each $0 < \varepsilon < 1/3$, the connected real-analytic closed curve of Theorem 5.1 is fixed once ε is chosen: its construction matches the jets of the three reference helices at the base phase and controls the approximation error on a phase window that shrinks with n, yielding $\Omega(n^{4/3-\varepsilon})$. Whether a single connected real-analytic closed curve can attain the exact exponent $4/3$ remains open.

Theorem 5.1. *For each $0 < \varepsilon < 1/3$, there exist a real-analytic embedded circle $\Gamma \subset \mathbb{R}^3$, two open balls U_-, U_+ with disjoint closures in $\operatorname{int}\operatorname{conv}\Gamma$, positive constants c, c_1, and an integer n_0. Draw a sequence of independent sites uniformly by arclength on Γ. For each $n \geq n_0$, on an event of probability at least $1 - \exp(-c_1 n^{1/3})$, the first n sites have at least $cn^{4/3-\varepsilon}$ Delaunay tetrahedra whose interiors meet* every *segment joining U_- to U_+. Almost surely this conclusion holds for all sufficiently large n. In particular, each such fixed segment s satisfies*

$$\liminf_{n\to\infty} \frac{\mathbb{E}N_n(s)}{n^{4/3-\varepsilon}} \geq c.$$

The curve admits a trigonometric polynomial parametrization, so each coordinate is a finite linear combination of 1, $\sin(ju)$ *and* $\cos(ju)$ *for positive integers j and angular parameter u. Almost surely no four sample sites are coplanar and no five are cospherical.*

Appendix C proves Theorem 5.1 in three steps: building Γ in Subsection C.1, locating the empty circumballs in Subsection C.2, and counting the resulting tetrahedra in Subsection C.3. Below we sketch the geometric mechanism behind those estimates and present two variants that eliminate the exponent loss ε.

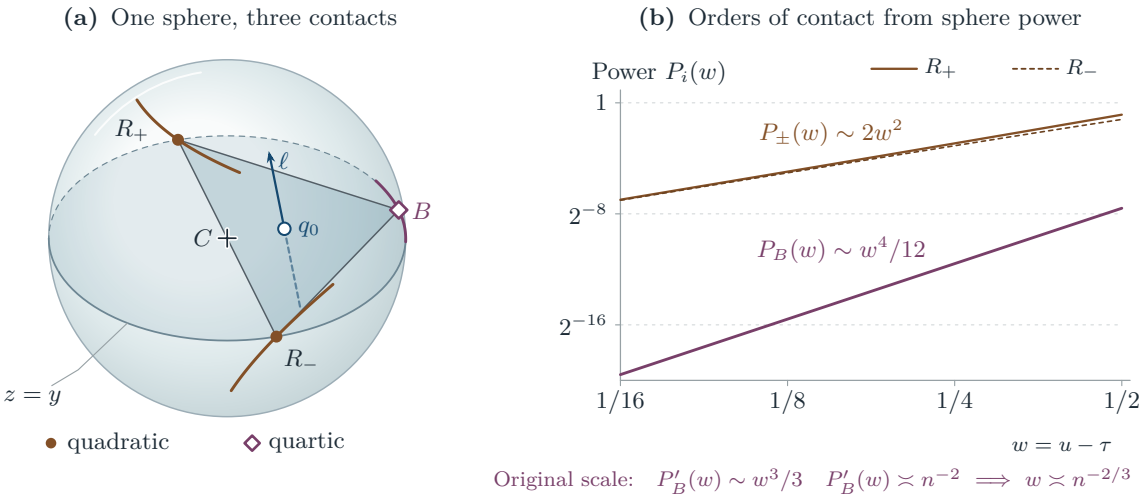


Figure 9: Geometry and power scaling of a fourth-order contact. In (a), the query line ℓ meets the interior of the contact triangle BR_+R_- at q_0, while the circle in the plane $z = y$ has the same tangent line and curvature vector as the front helix at B and also passes through R_+ and R_-. In (b), plotting sphere power against phase offset $w = u - \tau > 0$ on logarithmic axes contrasts the nearly coincident quadratic curves at $R_\pm$ with the much flatter quartic curve $w^4/12$ at B. Differentiating the front power $w^2 - 2(1 - \cos w)$ with respect to w gives $2(w - \sin w)$, asymptotic to $w^3/3$, which reaches the rear-perturbation scale n^{-2} when w is of order $n^{-2/3}$.

Where does the exponent 4/3 come from? Consider three ideal helical arcs generated by the screw motion H_u that rotates the horizontal plane through angle u about the vertical axis and translates vertically by u. Applying H_u to three base points B, R_+, R_- and to the centre C (the midpoint of R_+R_-) sweeps out a sphere of radius 2 at phase τ that tracks the helices while keeping all three contacts (Fig. 9(a)). Because H_u is an isometry, the power of the phase-τ sphere at a point of phase $u = \tau + w$ depends only on the offset w. It grows like $w^4/12$ along the front helix through B and like $2w^2$ along each rear helix through $R_\pm$ (Fig. 9(b)). At small offsets w, the front helix therefore hugs the sphere far more closely than either rear helix does.

Now fix one sampled site near R_+ and one near R_-, separated in phase by order $h = 1/n$, and join them into a rear edge. Any consecutive pair of front sites near B completes a tetrahedron thickening the reference triangle BR_+R_-. Every segment joining U_- to U_+ therefore pierces both faces incident to the rear edge and passes through the tetrahedron's interior. Why can the *same* rear edge support order $n^{1/3}$ empty Delaunay tetrahedra at once? Shifting the reference sphere by phase $O(1/n)$ alters its power at the two rear sites by order n^{-2}. Subtracting an affine function of slope $O(n^{-2})$ tilts the sphere to pass through both rear sites and one front site. Because the front power has slope asymptotic to $w^3/3$, that $O(n^{-2})$ tilt balances the front slope across an interval of width $w \asymp n^{-2/3}$, which holds $\Theta(n \cdot n^{-2/3}) = \Theta(n^{1/3})$ front sites on average. Suppose both rear sites sit inside a guard interval of width $2Dh = O(1/n)$ empty of other rear sites. Then Proposition C.1 shows that every consecutive front pair in this $n^{-2/3}$-window forms an empty circumball with that rear edge, generating a fan of $\Theta(n^{1/3})$ crossed tetrahedra (drawn for three cells in Fig. 15(b)).

To turn this local fan into a global lower bound, we count disjoint usable rear pairs along a closed curve. Theorem 5.1 closes the three helical arcs into one trigonometric polynomial loop matching position and the first M derivatives at each contact ($M \geq 4$). Choose M so that $\varepsilon(M - 1) > 2$. On the interval of width $n^{-\varepsilon}$, the errors in the power and its first two derivatives are then $O(n^{-\varepsilon(M-1)}) = o(n^{-2})$. This keeps the errors smaller than the n^{-2} perturbation used to tilt the sphere. That interval supplies order $n^{1-\varepsilon}$ disjoint rear guards and hence $\Omega(n^{4/3-\varepsilon})$ distinct tetrahedra, even though different front windows overlap. Taking

$\varepsilon = 1/12$, for example, yields $\Omega(n^{5/4})$, and a single M serves all $2/(M-1) < \varepsilon < 1/3$.

If instead we keep the three helical arcs exact over a fixed interval and join their six endpoints smoothly in $\mathbb{R}^3$, the usable interval no longer shrinks with n. We then gain order n disjoint rear pairs on a single C^∞ closed curve and recover the full exponent $4/3$.

Corollary 5.2. *There exist a C^∞ embedded circle $\Gamma_\infty \subset \mathbb{R}^3$, two open balls U_-, U_+ with disjoint closures contained in* $\operatorname{int}\operatorname{conv}\Gamma_\infty$*, positive constants c, c_1, and an integer n_0 with the following property. Sample a sequence of independent sites uniformly by arclength on Γ_∞. For each $n \geq n_0$, with probability at least $1 - \exp(-c_1 n^{1/3})$, at least $cn^{4/3}$ tetrahedra belong to every Delaunay triangulation of the first n sites and each of their interiors meets every segment joining U_- to U_+. Almost surely this holds for all sufficiently large n. In particular, each fixed joining segment s satisfies* $\liminf_{n\to\infty} n^{-4/3}\mathbb{E}N_n(s) \geq c$.

As proved on page 52, Corollary 5.2 follows directly from the exact-family estimate of Appendix D (Proposition D.2). Because the screw motion keeps all three contact orders exact over a fixed interval, no derivative-matching approximation is needed and the full $n^{4/3}$ rate holds.

Finally, a second construction achieves exact contact orders over a fixed interval without losing real-analyticity: place the front arc and the two rear arcs on two separate real-analytic closed curves. Via stereographic projection from $S^3 \subset \mathbb{R}^4$, Corollary 5.3 builds two disjoint analytic circles (Fig. 16 in Appendix D) that again feed into Proposition D.2 to reach the full $n^{4/3}$ lower bound.

Corollary 5.3. *There exist two disjoint real-analytic embedded circles in $\mathbb{R}^3$ whose union Γ has the following property. There are open balls U_-, U_+ with disjoint closures in* $\operatorname{int}\operatorname{conv}\Gamma$*, positive constants c, c_1, and an integer n_0. Draw a sequence of independent sites from normalized arclength on Γ. The probability of sampling either component is its fraction of the total length. For each $n \geq n_0$, with probability at least $1 - \exp(-c_1 n^{1/3})$, the first n sites have at least $cn^{4/3}$ Delaunay tetrahedra whose interiors meet every segment joining U_- to U_+. This holds eventually almost surely. In particular, every fixed joining segment s satisfies* $\liminf_{n\to\infty} n^{-4/3}\mathbb{E}N_n(s) \geq c$.

Subsection D.2 proves Corollary 5.3 by pulling back explicit trigonometric factorizations through stereographic projection from S^3 and applying Proposition D.2. It also verifies almost sure general position of samples drawn across both components so that every empty four-site circumball forms a valid Delaunay tetrahedron.

6 Computational experiments on Delaunay walks

Theorem 1.1 bounds the expected crossing count asymptotically in both the sample size n and the segment length $|s|$, while Sections 4 and 5 show how contact degeneracies can raise the growth exponent. How clearly do these regimes emerge at finite sample sizes? To check, we ran three complementary experiments. We first varied n with s fixed to measure empirical exponents across intrinsic and ambient dimensions. We next inspected individual tetrahedra on two space curves to see how often the walk uses the four-site pattern of Corollary 5.2. Finally, we quadrupled $|s|$ on fixed triangulations to isolate the role of segment length and contact motion.

Within each comparison, the same query segment coordinates were evaluated across all point sets at a given sample size. In every series except the follow-up pole-axis check on F in Subsection 6.3, the query segment was fixed before generating the samples; for that pole-axis check, the three axial segments were chosen after inspecting the perpendicular direction on the same 32 point clouds, with their endpoints fixed before computing any counts along the pole

axis. Samples were drawn independently and uniformly by volume in solid bodies, by intrinsic volume on manifolds, and by arclength on curves. Figure 10 counts internal facets crossed (each step from one simplex into an adjacent one), whereas Figs. 11 and 12 count visited cells, including the initial cell when the segment enters the sample hull. Where shown, pointwise 95% bootstrap confidence intervals (CIs) were computed by resampling entire point sets with replacement; for paired length comparisons, each resampled point set contributed both its short and long counts.

Our method for extracting the crossed cells depended on the ambient dimension. On space curves in $\mathbb{R}^3$, we built exact 3D Delaunay triangulations with CGAL 6.1.1 [23], where intersection predicates and crossing points are exact for the stored coordinates. In all 640 curve samples, tracing the segment backwards reproduced the exact same sequence of tetrahedra in reverse order. For manifolds in $\mathbb{R}^4$ and $\mathbb{R}^5$, a Python routine located the crossed simplices directly via the paraboloid lift $(p, |p|^2)$. It called the HiGHS linear solver [21] to identify a starting cell, tracked entry and exit intervals in exact rational arithmetic, and verified that every reported simplex had an empty open circumball against all sampled sites. For the surface G in $\mathbb{R}^6$, we used exact predicates on the stored coordinates of all 128 point sets. We selected one point set in advance at each of the two extreme sizes, $n = 1{,}024$ and $n = 65{,}536$, for two further checks. Reversing the segment reproduced the cells and the interval endpoints in reverse order. We also regenerated the coordinates from the same random proposals using 84 decimal digits and rounded them to a grid of step 2^{-120}. The accepted indices and the forward sequence of cells agreed. A second point set at $n = 1{,}024$, included in the same 128, also gave the same sequence under reversal. The solid-ball, ellipsoid and torus baselines reuse earlier counts obtained by verified linear programming or exact lifted-hull computation. The accompanying research data and software document the sampling and traversal methods and provide the observations used in the comparisons below.

6.1 How does the count grow with the number of sites?

To separate the effect of intrinsic dimension δ from that of ambient dimension d, we varied n from $1{,}024 = 2^{10}$ to $65{,}536 = 2^{16}$ with the segment fixed. Under this 64-fold increase in sample size, the scale $n^{1/\delta}$ grows by a factor of 8 when $\delta = 2$ and by a factor of 4 when $\delta = 3$. Write $M(n)$ for the measured mean number of crossed internal facets. We track both the growth factor $M(2^{16})/M(2^{10})$ and the log–log slope $\widehat{\beta} = \log[M(2^{16})/M(2^{10})]/\log 64$, comparing them with the upper-bound exponent $1/\delta$ (which also allows slower growth).

We tested embedded manifolds whose supporting geometry stays simple along the entire walk. At each point q on the segment, a unique supporting ball touches the manifold at two moving ordinary contacts whose chord has midpoint q. Starting from the profile $(x, u) \in \mathbb{R} \times \mathbb{R}^\delta$ defined by $x^2 + |u|^2 + \kappa|u|^4/2 = 1$, we took the quadratic profile ($\kappa = 0$) in $\delta \in \{2, 3\}$ and the quartic profile ($\kappa = 1$) in $\delta = 2$. We embedded each profile into both $\mathbb{R}^4$ and $\mathbb{R}^5$ via $H(x, u) = (x, 2u, xPu)$, where $P\colon \mathbb{R}^\delta \to \mathbb{R}^{d-\delta-1}$ is a linear projection with orthonormal rows (omitted when $d = \delta + 1$), and traced the same segment of length $L_0 = \sqrt{5}/4$ in the shared coordinates. Appendix A.6 gives the explicit maps and verifies regularity along this segment.

On the quadratic profile in $\mathbb{R}^4$, multiplying n by 64 increased the mean facet count by a factor of 8.28 for $\delta = 2$ and 4.28 for $\delta = 3$. In $\mathbb{R}^5$, the corresponding factors were 7.44 and 3.86 (Fig. 10(b)). Across both ambient dimensions, raising δ from 2 to 3 reduced the measured exponent $\widehat{\beta}$ by roughly 0.16, matching the theoretical drop $1/2 - 1/3 \approx 0.167$. By contrast, raising the ambient dimension from $d = 4$ to $d = 5$ at fixed profile shifted $\widehat{\beta}$ by only -0.026 and -0.025 for the quadratic profiles and -0.002 for the quartic profile (Fig. 10(a) displays the individual slope estimates and their pointwise 95% bootstrap CIs). For these contrasts (see `dimension_scaling/slope_contrasts.csv`), all three pointwise 95% bootstrap intervals for changing d alone include zero ($[-0.0516, 0.0006]$, $[-0.0628, 0.0126]$ and $[-0.0239, 0.0214]$),

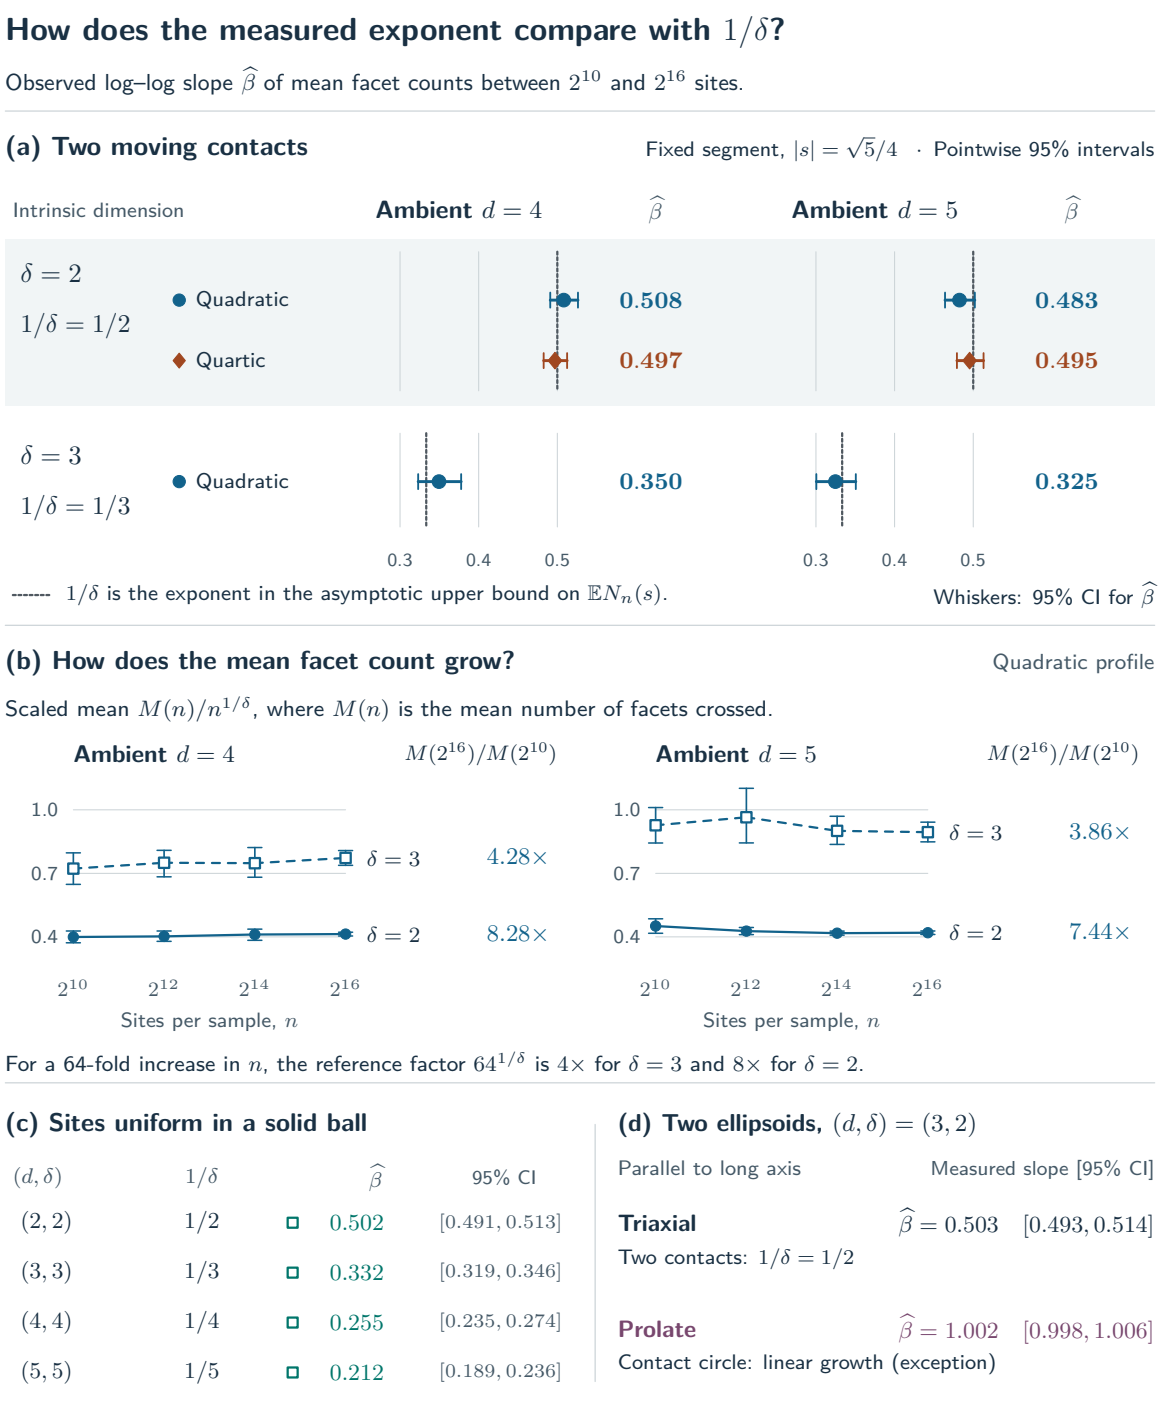


Figure 10: Empirical growth of the mean number of internal facets crossed compared with the theoretical upper-bound exponent $1/\delta$. In (a), moving down a column changes the intrinsic dimension δ at fixed ambient dimension d, while moving across a row changes d for the same profile; dashed vertical lines mark the theoretical bound exponent $1/\delta$, and symbols give the individual measured endpoint slopes $\widehat{\beta}$ between 2^{10} and 2^{16} sites (32 independent samples at each endpoint size, and 16 at intermediate sizes except for the quadratic surface in $(d, \delta) = (5, 2)$, which uses 32 at every size; slope-difference CIs across d appear in the text). In (b), normalized means $M(n)/n^{1/\delta}$ remain nearly level across the four tested sizes, consistent with $n^{1/\delta}$ scaling over this range; note that the $\delta = 2$ and $\delta = 3$ curves use different divisors $n^{1/\delta}$ and cannot be compared vertically as raw counts, while numbers at right give unscaled growth ratios. Solid balls of segment length 0.6 appear in (c), and two ellipsoids in $\mathbb{R}^3$ along a segment of length 1.2 parallel to the long axis appear in (d). Whiskers and intervals show pointwise 95% bootstrap CIs for the individual estimates.

which is compatible with no shift in finite-sample slope (though it does not prove equality of asymptotic exponents). Throughout the intermediate sizes in Fig. 10(b), dividing $M(n)$ by $n^{1/\delta}$ yields nearly level curves, showing that the finite-sample growth over the tested range closely tracks $n^{1/\delta}$ while the multiplicative constant adapts to the embedding. As a baseline check for the solid case $\delta = d$ of Proposition 2.1, solid balls in dimensions two through five yielded slopes $0.502, 0.332, 0.255$ and 0.212 (Fig. 10(c)), matching $1/2, 1/3, 1/4$ and $1/5$.

At the same time, Fig. 10(d) and our higher-dimensional switching example confirm that dimension alone cannot determine the exponent when contact geometry changes. On two ellipsoids in $\mathbb{R}^3$, following a segment of length 1.2 parallel to the long axis gave slope 0.503 on a triaxial ellipsoid (two isolated ordinary contacts, Proposition 4.4) against 1.002 on a prolate spheroid (circular contact slices advancing along the segment, Corollary 4.3). Conversely, on the surface parametrized by G in $\mathbb{R}^6$ from Subsection 3.3, three contacts remain stationary on each half of the segment and only one contact switches at the midpoint. Comparing the endpoint sizes 2^{10} and 2^{16}, its mean facet count scarcely changed, moving from 3.53 to 3.56 (with intermediate means 4.03 and 3.03 at 2^{12} and 2^{14}).

6.2 How often do the crossed tetrahedra have the four-site pattern?

Corollary 5.2 proves its $n^{4/3}$ lower bound by counting only tetrahedra of a specific combinatorial type: one vertex on each rear arc forming a shared rear edge r_+r_-, and two consecutive sites p_i, p_{i+1} on the front arc near the fourth-order contact (Fig. 11(a)). Sampling sites independently and uniformly by arclength on each curve and tracing the same fixed vertical segment of length $1/16$, we tested how dominant this pattern is on the smooth curve Γ_∞. We also compared it against an explicit closed real-analytic curve Γ_5 passing through the same three reference points (parametrized in Appendix A.6).

On the smooth curve Γ_∞, the pooled fraction of visited tetrahedra matching the four-site pattern (total marked tetrahedra divided by total crossed tetrahedra summed across all walks at each size) climbed steadily with n, rising from 48% at $n = 256$ to 72% at $n = 1{,}024$ and 81% at $n = 4{,}096$ (Fig. 11(b)). At $n = 4{,}096$, each active rear edge contributed an average of 5.37 marked tetrahedra within the same walk, directly reflecting the fan of consecutive front pairs anticipated by the proof.

Comparing Γ_∞ with the analytic curve Γ_5 reveals an instructive twist (Fig. 11(c)). At $n = 4{,}096$, walks on Γ_5 crossed *more* tetrahedra on average than walks on Γ_∞ (442 versus 277), yet fewer of them exhibited the three-arc four-site pattern (99 versus 225). Across the common experimental phase window $|\tau| \le 1/32$, the two-curve contact analysis in Appendix A.6 (page 44) explains both contact counts: on the sampled smooth curve Γ_∞, three exact helical plateaus keep $R = 2$, the smooth radial bump lifts the fourth plateau outside the moving sphere, and phase separation keeps the connecting transitions strictly outside the sphere, leaving exactly three contacts; on Γ_5, keeping $R \equiv 2$ gives four roots of $\sin(2v) = 4\tau$ and hence four contacts (including $(-3, 0, 0)$ at $\tau = 0$). The walk on Γ_5 was longer because it crossed more tetrahedra with other vertex arrangements; determining how many of those tetrahedra use sites near the fourth contact would require classifying their vertices as well.

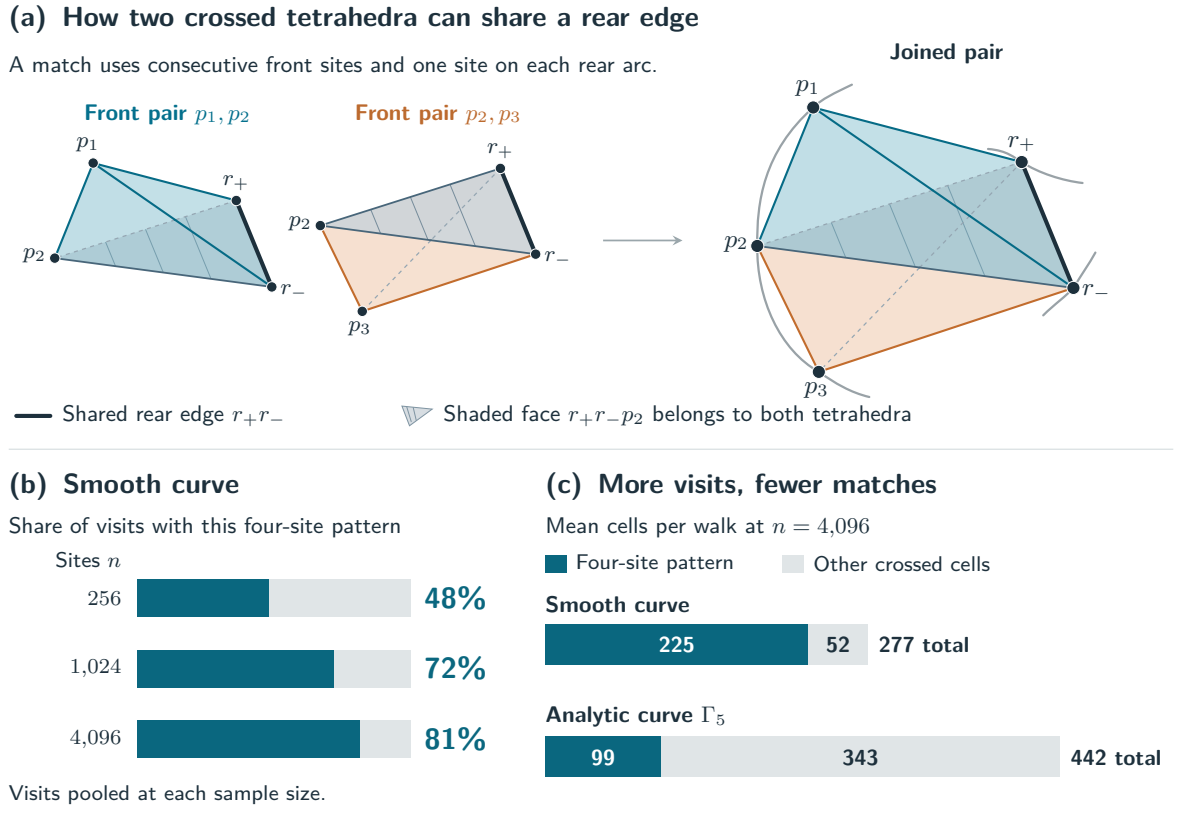


Figure 11: Prevalence of the four-site pattern on two space curves along a fixed vertical segment of length 1/16. In (a), consecutive front sites p_1, p_2, p_3 combine with one fixed rear pair r_+r_- to form adjacent tetrahedra sharing the shaded face $r_+r_-p_2$. In (b), pooling 256, 32 and 32 walks on the smooth curve Γ_∞ at $n = 256, 1{,}024$ and $4{,}096$ shows the ratio of marked to total crossed tetrahedra rising from 48% to 81%. In (c), mean counts per walk at $n = 4{,}096$ (32 samples each) show why total counts alone can mislead: Γ_5 crosses more tetrahedra overall (442 vs. 277), yet fewer match the three-arc pattern (99 vs. 225), leaving 343 unclassified cells per walk on average (across $|\tau| \le 1/32$, each moving sphere touches Γ_∞ at three points and Γ_5 at four, while our classifier marks only the three-arc pattern).

6.3 By what factor does the cell count change when the length quadruples?

Our upper bound $O(1 + |s|n^{1/\delta})$ scales linearly with $|s|$ once the walk leaves the first cell, yet leaves open whether a longer segment actually incurs proportional crossings along every direction. To isolate the effect of segment length, we reused 32 samples of 65,536 sites on the quadratic surface $F(x, u) = (x, 2u, xu)$ in $\mathbb{R}^5$ ($\delta = 2$) from Subsection 3.3. On each sample, we tested three concentric segments of lengths $L_0/4, L_0/2, L_0$ ($L_0 = \sqrt{5}/4$) along two perpendicular directions through the origin (Fig. 12(a)). We first traced the segments in direction $(0, 2, 1, 0, 0)$, perpendicular to the axis joining the poles $(\pm 1, 0, 0, 0, 0)$. Afterwards, we chose the three segments along the pole axis and traced them on the same triangulations, fixing their endpoints before computing any counts along that axis.

Along the pole axis, the unit sphere touches $F(S^2)$ solely at the two poles and supports every interior point between them. On a compact regular subset, the common-ball refinement of Theorem 3.2 (Appendix A.1) yields an $O(1)$ upper bound on the expected count independent of n. In our 32 finite samples at $n = 65{,}536$, even the longest pole-axis segment lay inside a single Delaunay simplex in every trial (we verified in each cloud that both endpoints belonged to that simplex and no other site fell inside its circumsphere), giving an observed long-to-short ratio of 1.00. Along the perpendicular line in direction $(0, 2, 1, 0, 0)$, Appendix A.6 gives the two ordinary contacts and the supporting sphere as smooth functions of the query point. Along

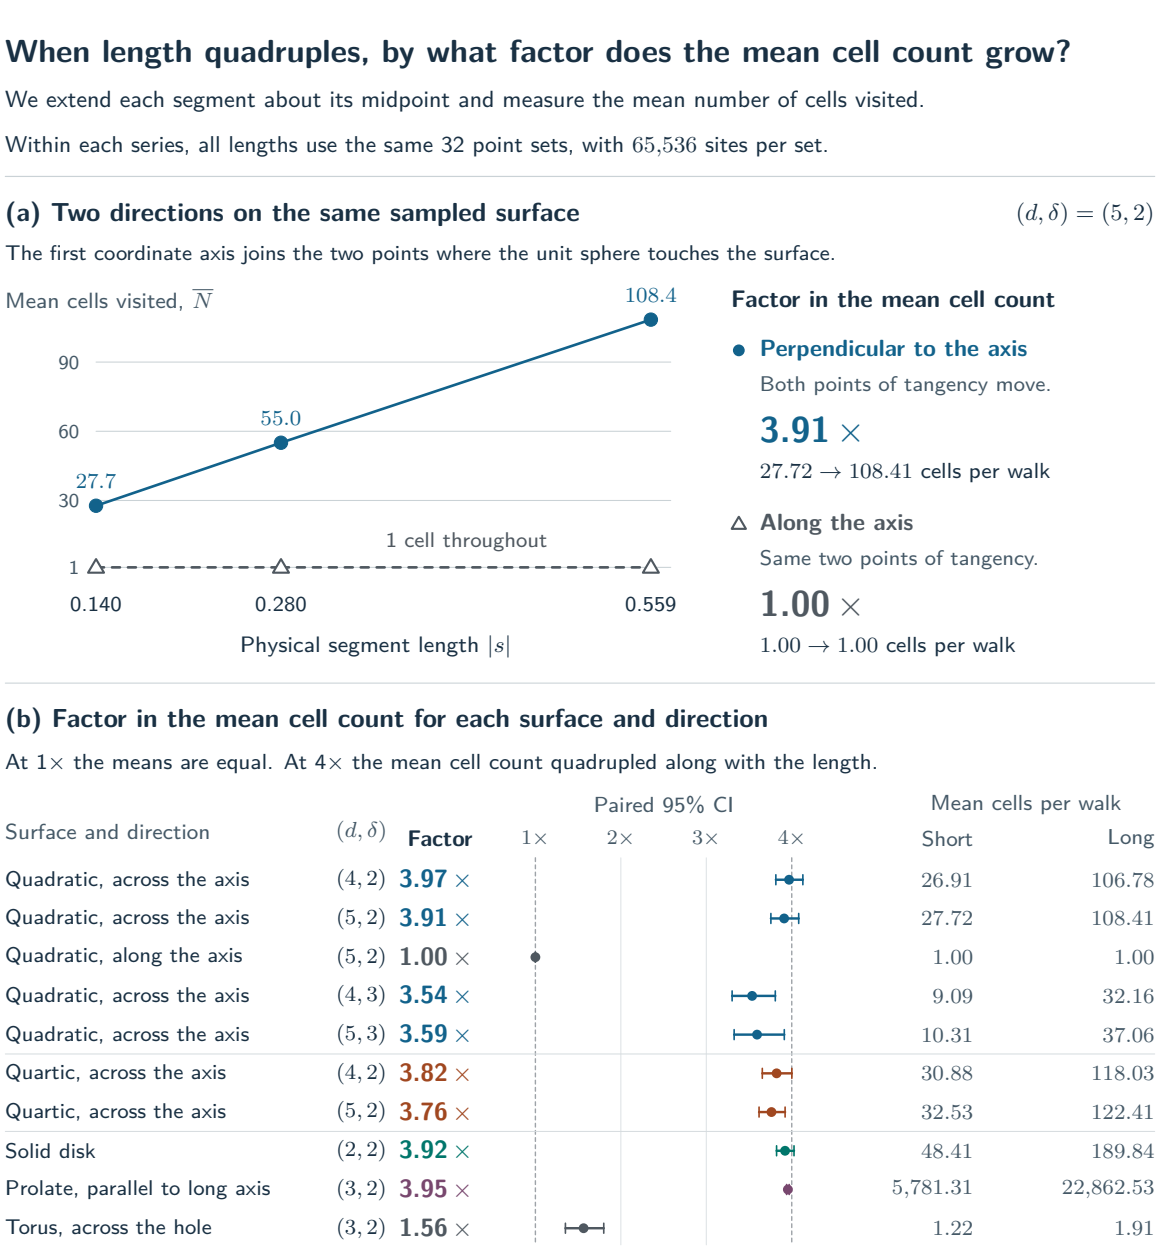


Figure 12: Effect of quadrupling segment length about its midpoint at fixed sample size $n = 65{,}536$ (32 independent point sets per row, preserving pairing for the 95% bootstrap CIs). In (a), two perpendicular lines on the same 32 samples of $F(S^2) \subset \mathbb{R}^5$ show opposite finite-sample scaling: along the pole chord (supported on the continuous surface by one fixed sphere), all three segment lengths stayed inside a single cell in every sample (1.00×), whereas along the perpendicular direction, where the two contacts move, quadrupling length multiplied the mean by 3.91× (95% CI $[3.76, 4.08]$). In (b), eight series show paired factors between 3.54× and 3.97× arising from three distinct mechanisms: the six regular moving-contact manifold profiles (Theorem 3.2), the solid disk under volume sampling (Proposition 2.1), and the prolate spheroid walk parallel to the long axis across circular contact slices (Corollary 4.3, already exceeding 5,700 cells); by contrast, the pole axis of F stays at 1.00× and the central-hole torus segment stays inside one inner contact disk and grows by only 1.56×. Shorter lengths are $\sqrt{5}/16$ for the quadratic and quartic surfaces, 0.15 for the solid disk and torus, and 0.30 for the prolate spheroid.

this direction, where the contacts move across the surface, the observed mean cell count rose from 27.72 to 55.03 and 108.41, giving a paired long-to-short ratio of 3.91 (95% bootstrap CI $[3.76, 4.08]$).

Figure 12(b) places this directional contrast across all ten length-comparison series at

$n = 65{,}536$. Eight series exhibit similar finite-sample factors between 3.54 and 3.97 when $|s|$ quadruples, yet they arise from three distinct theoretical regimes: the six regular moving-contact profiles in $\mathbb{R}^4$ and $\mathbb{R}^5$ (governed by the upper bound of Theorem 3.2), the solid planar disk under volume sampling (Proposition 2.1), and the segment parallel to the long axis of the prolate spheroid at constant transverse offset 0.071, where circular contact slices drive linear growth in n (Corollary 4.3). Indeed, on the prolate spheroid that linear axial scaling already pushes the short-segment count above 5,700 cells, and extending the segment fourfold multiplies it by 3.95 to over 22,800 cells, showing that similar length ratios at a single fixed n do not by themselves distinguish regular $n^{1/\delta}$ scaling from degenerate $\Theta(n)$ growth. By contrast, on a torus of major radius 2 and tube radius 1/2, a horizontal segment through $(0.071, 0, 0.053)$ in the central hole stays inside the contact disk of one inner latitude circle (a non-regular single-sphere geometry, distinct from the regular two-pole chord on F), yielding a modest factor of 1.56 (from 1.22 to 1.91 cells). Together, Figs. 12(a) and (b) illustrate at fixed n how stationary and moving supports separate across the tested series, while distinguishing regular moving-contact profiles from volume and circular-contact regimes and leaving general geometric conditions for a matching $\Omega(|s|n^{1/\delta})$ lower bound as an open problem.

7 Conclusion

At first sight, a straight Delaunay walk across the interior of a sampled manifold poses a puzzle. The segment moves through empty ambient space where there are no sites, yet the spacing $n^{-1/\delta}$ along the manifold still governs how many cells it crosses. This paper identifies the geometric mechanism behind that scaling. Under the paraboloid lift, every interior query point is supported by empty spheres tangent to the manifold, and our analysis shows that nondegenerate (ordinary) contacts whose tangent spaces and chords span $\mathbb{R}^d$, together with generic codimension-one switches across the divisible interfaces $d = j(\delta + 1)$, suffice to keep the expected cell count small. For an embedding in the residual C^∞ class of Theorem 1.1 (for a compact smooth δ-manifold without boundary and of full affine span in $\mathbb{R}^d$, $1 \le \delta < d$), almost every pair of interior endpoints defines a fixed segment s with $\mathbb{E}N_n(s) \le C_s(1 + |s|n^{1/\delta})$, and the constants stay uniform over all fixed segments contained in a compact subset $Q \subset G$. Likewise, for a specific C^2 embedding, Theorem 3.2 yields uniform constants over every fixed segment in a compact subset $Q \subset \mathcal{R}$ of regular queries. In parallel, Proposition 2.1 settles the solid-body baseline ($\delta = d$) for a fixed sample size n in arbitrary dimension $d \ge 1$, removing the logarithmic factor of the three-dimensional walk bound [27] down to the natural boundary margin $An^{-1/d}$.

Stepping back, these upper bounds separate total mesh complexity from the local cost of a single query. Earlier work on surfaces in $\mathbb{R}^3$ focused mainly on the *total* number of cells in the triangulation, under deterministic (ε, κ)-sampling on C^5 surfaces satisfying their genericity hypotheses [4] or under random Poisson sampling [12, 17]. By contrast, a straight walk probes only a one-dimensional slice controlled by local supporting balls along s. Within a compact regular region $Q \subset \mathcal{R}$ of Theorem 3.2 (where contacts are nondegenerate and their tangent spaces and chords span $\mathbb{R}^d$), if a single ball supports every point of a fixed segment $s \subset Q$, the expected crossing count stays $O(1)$ as n grows. When supporting contacts move along a fixed segment s that stays inside a compact subset $Q \subset G$ (for an embedding in the residual class of Theorem 1.1) or inside a compact subset $Q \subset \mathcal{R}$ (for a C^2 embedding under Theorem 3.2), the expected count remains bounded above by $O_Q(1 + |s|n^{1/\delta})$. Crucially, moving contacts by themselves do not guarantee this rate: on the prolate spheroid in $\mathbb{R}^3$ ($\delta = 2$), the supporting contacts move along an axial segment yet yield $\Theta(n)$ expected crossings (Corollary 4.3). Our finite-sample length comparisons (Fig. 12) reinforce both lessons, contrasting stationary and moving contacts along two directions on the same surface while showing that similar near-

fourfold length ratios can arise from regular moving-contact profiles, volume sampling, or circular-contact axial walks.

Our two counterexample families reveal different limits of the contact hypotheses and show why intrinsic dimension alone cannot replace generic position. First, on a prolate spheroid in $\mathbb{R}^3$ (Section 4), each interior supporting sphere touches the surface along an entire circular latitude whose convex hull forms a planar contact disk. A fixed segment confined to one contact disk can therefore have bounded expected cost $O(1)$ even though its supporting sphere touches a full circle, showing that isolated contacts are not necessary for a bounded walk on a particular segment; by contrast, a fixed interior segment with nonzero axial displacement sweeps across a continuum of contact disks and visits $\Theta(n)$ cells in expectation, exceeding $n^{1/2}$. The same axial dichotomy extends to higher-dimensional spheroids (where each slice is a $(d-2)$-sphere) and to the analytic products $\mathcal{E} \times M$ of Corollary 4.2, where each interior contact slice consists of two copies of M (for instance, two 2-spheres when $d = 5$ and $\delta = 3$): the common property used by the proof is that the convex hull of each contact slice fills the entire cross-section at that axial level. Second, isolated contacts alone do not suffice for the intrinsic bound. In $\mathbb{R}^3$, both the smooth closed curve of Corollary 5.2 and the two-component real-analytic link of Corollary 5.3 sustain an *exact* one-parameter family of empty spheres with only three isolated contacts, one of fourth order and two quadratic, over a *fixed* parameter interval (Proposition D.2). Under uniform arclength sampling on the smooth curve, or on the union of the two real-analytic components with probabilities proportional to their lengths, with probability tending to one, at least $cn^{4/3}$ Delaunay tetrahedra simultaneously meet every segment joining two fixed interior balls, giving expected count $\Omega(n^{4/3})$ above the intrinsic scale n for each fixed joining segment. By contrast, for each $0 < \varepsilon < 1/3$, our connected real-analytic closed curve in $\mathbb{R}^3$ (Theorem 5.1) approximates that contact geometry by jet matching over a *shrinking* phase window of width $n^{-\varepsilon}$, yielding a fixed curve chosen as a function of ε whose corresponding simultaneous lower bound is $\Omega(n^{4/3-\varepsilon})$. Attaining the exact exponent $4/3$ on a single connected real-analytic closed curve remains open.

Five open problems emerge naturally from this contact-geometric picture and invite further work across computational geometry, probability and real algebraic geometry. First, our upper bounds and directional experiments expose a sharp contrast between stationary and moving contacts without settling the converse. Pinning down the geometric conditions under which moving ordinary contacts along a fixed segment $s \subset Q \subset G$ (or $s \subset Q \subset \mathcal{R}$) force cell boundaries to split consecutive pieces of length $n^{-1/\delta}$ would establish when a matching $\Omega(|s|n^{1/\delta})$ lower bound accompanies Theorems 1.1 and 3.2. Second, in Theorem 5.1 the connected real-analytic closed curve in $\mathbb{R}^3$ depends on ε and loses a factor n^{ε} relative to the smooth and two-component constructions. Can a *single* connected real-analytic closed curve in $\mathbb{R}^3$ realize the exact contact family of Proposition D.2 over a fixed interval to attain $\Omega(n^{4/3})$ without an ε loss? Closely related, since the proof of Theorem 1.1 shows that almost every interior segment avoids the exceptional set of queries associated with degenerate contacts whereas the constructions in Corollaries 5.2 and 5.3 use a persistent fourth-order contact, which higher-order isolated contacts can a walk cross while still obeying $O(1 + |s|n^{1/\delta})$?

Two distinct boundary questions also remain open. Third, for a compact smooth manifold without boundary embedded in the residual class of Theorem 1.1 (recall that Sections 4 and 5 already rule out uniform $n^{1/\delta}$ scaling without genericity), can the constants be chosen uniformly over all fixed segments contained in the good open set G, including segments approaching the convex-hull boundary ∂K arbitrarily closely, thereby moving toward the full strength of [11, Conjecture 2]? Fourth, how does the contact analysis extend to compact manifolds *with* boundary (such as trimmed CAD surfaces or the finite lateral surface of a circular cylinder, bounded by two boundary circles), and separately to topologically closed *piecewise-smooth* surfaces with sharp edges or corners (such as the closed circular cylinder with end caps tested in [11]), where supporting contacts can migrate onto boundary curves, seams, or corners and

cease to satisfy interior tangency equations?

Fifth, our upper bounds concern segments fixed before drawing the sample, whereas the simultaneous high-probability events of Theorem 5.1 and Corollaries 5.2 and 5.3 already allow choosing any segment joining U_- to U_+ after observing the sample. Determining the expected *worst-case* crossing number over all segments or lines chosen after sampling on a generic manifold remains wide open, calling for a manifold analogue of the planar and cubic stabbing bounds in [7, 28]. Equally open is whether the randomized or visibility-based walks used in geometric software can sidestep axial contact circles and fourth-order fans altogether. Taken together, our theorems and counterexamples show that point location in manifold-sampled Delaunay triangulations is governed by an interplay between intrinsic sampling density and the contact geometry of empty spheres. Understanding that interplay is what turns an ambient mesh into a predictable geometric search structure.

A Proof of the generic crossing bound

Theorem 1.1 asserts an upper bound along an entire segment $[a, b] \subset \operatorname{int} K$, which can traverse regular queries as well as codimension-one interfaces where supporting contacts change. Our proof runs in five subsections. Subsection A.1 handles compact subsets of the regular set $\mathcal{R}$, proving Theorem 3.2. Subsection A.2 uses jet transversality to bound the dimension of all degenerate and deficient sets, and Subsection A.3 treats queries right on a sampled hypersurface ($\delta = d - 1$). Subsection A.4 then analyses the closed interval of supporting centres at the divisible interfaces $d = j(\delta + 1)$, and Subsection A.5 pairs both endpoint spheres to prove Proposition A.1 across the interface. Combining these pieces completes Theorem 1.1 below, after which Subsection A.6 verifies the contact geometry behind the experiments of Section 6.

A.1 From spanning contacts to a local count

To prove Theorem 3.2, fix a compact set $Q \subset \mathcal{R}$ of regular queries. First we show that the supporting balls over Q form a compact family and that $\mathcal{R}$ is open. Next, for each $q \in Q$ and scale $t > 0$, we place 2δ signed witness patches around each contact in a minimal representation so that picking one site per patch encloses a uniform ball $B(q, \gamma t)$, and we add one patch around each continued local minimum of power. Lemma 3.3(b) then forces all vertices of any simplex reaching $B(q, \gamma t/2)$ into those $O(t)$-balls. Lemma 3.1 bounds the local expected cell count, and a covering plus Fubini's projection argument yields the segment bound.

Proof of Theorem 3.2. First check that the family of pairs (q, B), where $q \in Q$ and $B = B(c, R)$ is a supporting ball at q, is compact. Each such ball defines via (2.3) an affine support $\ell(x) = \ell(q) + 2c \cdot (x - q)$ with $\ell(q) = L_V(q)$ and $\ell(x) \le L_V(x)$ on K. Because Q is compact inside $\operatorname{int} K$, some fixed radius $r_0 > 0$ satisfies $B(q, r_0) \subset K$ for all $q \in Q$. Testing $\ell(x) \le L_V(x)$ at $x = q + r_0 c/|c|$ bounds $2r_0|c| \le \max_K L_V - \min_K L_V$, so c stays bounded, and contact with compact V bounds R as well. Given any convergent sequence $(q_k, c_k, R_k) \to (q, c, R)$, the open interior of the limiting ball still misses V. Extracting a convergent subsequence of Carathéodory representations (expressing each q_k as a convex combination of at most $d + 1$ contacts of B_k) places q in the convex hull of contacts of the limiting ball, which therefore supports q.

Why regular queries form an open set. Let $q \in \mathcal{R}$. Because V is compact and $q \notin V$, every sufficiently close point q' also misses V. Suppose for contradiction that points $q_k \to q$ failed regularity. If q_k had a supporting ball B_k with a degenerate contact $p_k \in V$, compactness above would extract a convergent subsequence to a supporting ball B at q touching V at a limiting contact p. The Hessian of the power restricted to V would then inherit a zero

eigenvalue at p, contradicting regularity of q. Otherwise, each q_k admits a supporting ball B_k and a minimal representation $q_k = \sum_{i=1}^{j} \lambda_{k,i} p_{k,i}$ whose contact tangents and chords are orthogonal to some unit vector $v_k \in \mathbb{R}^d$. Passing to a subsequence with fixed size j and limits $B_k \to B$, $p_{k,i} \to p_i$, $\lambda_{k,i} \to \lambda_i$, and $v_k \to v$, every contact p_i with limiting weight $\lambda_i > 0$ stays a contact of B satisfying $(p_i - q) \cdot v = 0$ and $T_{p_i} V \perp v$. Merging equal limit points and dropping zero-weight terms writes q as a positive convex combination of contacts of B; extracting a minimal subset gives a minimal representation at q orthogonal to v, again contradicting $q \in \mathcal{R}$. Hence $\mathcal{R}$ is open.

Enclosing the query. Fix $q \in Q$ and a supporting ball $B = B(c, R)$ at q, and write $q = \sum_{i=1}^{m} \lambda_i p_i$ as a minimal representation by contacts of B, with $\lambda_i > 0$ and $\sum_i \lambda_i = 1$. Near each contact p_i, write V locally as a C^2 graph over $T_{p_i} V$ centred at p_i, and choose an orthonormal basis $e_1, \ldots, e_\delta$ of $T_{p_i} V$. At small scale $t > 0$, place 2δ witness patches on V around p_i as the graph images of tangent balls of radius εt centred at $\pm t e_j$ ($1 \le j \le \delta$). For a small fixed $\varepsilon > 0$, each patch has sampling mass at least $a t^\delta$ by (1.1). Moreover, tangency at p_i keeps every patch point within distance $O(t)$ of p_i with power $g_B \le A t^2$. We claim that for a uniform constant $\gamma > 0$ independent of $q \in Q$, B and small t, picking one site from each of these $2\delta m$ witness patches forces their convex hull to contain $B(q, \gamma t)$.

To verify this inclusion, it suffices by convex separation to show that for every unit direction $v \in \mathbb{R}^d$, at least one chosen site x satisfies $v \cdot (x - q) \ge \gamma t$. Because we place patches in both directions $\pm t e_j$ at every contact p_i regardless of how small λ_i is, some signed basis vector at p_i has inner product at least $t |\operatorname{proj}_{T_{p_i} V} v| / \sqrt{\delta}$ with v. Moving within the tangent ball of radius εt alters the projection by at most $\varepsilon t |\operatorname{proj}_{T_{p_i} V} v|$, and normal bending of the graph costs at most $C t^2$. Fixing $\varepsilon < 1/(2\sqrt{\delta})$ yields a constant $b > 0$ such that the largest displacement from q in direction v among witness sites near p_i is at least

$$v \cdot (p_i - q) + bt \, |\operatorname{proj}_{T_{p_i} V} v| - C t^2.$$

Suppose for contradiction that no uniform $\gamma > 0$ works at arbitrarily small scales. Then there exist queries $q_k \in Q$, supporting balls B_k, scales $t_k \to 0$, and unit vectors v_k making this lower bound at most $o(t_k)$ for every i. Multiplying by $\lambda_{k,i} > 0$, summing over i, and dividing by t_k cancels the chord terms (since $\sum_i \lambda_{k,i}(p_{k,i} - q_k) = 0$) and leaves

$$b \sum_i \lambda_{k,i} |\operatorname{proj}_{T_{p_{k,i}} V} v_k| \le C t_k + o(1).$$

Extracting a convergent subsequence of queries, balls, contacts, weights and unit vectors $v_k \to v$ forces $v \perp T_{p_i} V$ at every contact with positive limiting weight $\lambda_i > 0$. Returning to the original per-contact inequality, dropping the nonnegative tangential term and letting $t_k \to 0$ directly yields $v \cdot (p_i - q) \le 0$ for all those contacts; because their λ_i-weighted average vanishes, every term with $\lambda_i > 0$ vanishes individually. Grouping coinciding positive-weight limits and extracting a minimal subset produces a minimal representation at the limit query $q \in Q$ orthogonal to v, contradicting $q \in \mathcal{R}$. The exact same argument applies inside a hyperplane when sites are projected orthogonally onto it and the contact directions span it, a fact we use in Subsection A.5.

Confining the vertices. For any fixed supporting ball B_0 at $q_0 \in Q$, every contact is ordinary, so contacts are isolated in compact V and hence finite in number. Around each contact of B_0, the implicit function theorem continues the local minimum of g_B to a unique local minimum $z_i = z_i(B)$ depending C^1 on the ball parameters for all balls B near B_0, obeying a uniform quadratic bound

$$a_i + c|p - z_i|^2 \le g_B(p) \le a_i + C|p - z_i|^2, \qquad a_i = g_B(z_i).$$

For every supporting ball B in this neighbourhood, $g_B \geq 0$ on V forces $a_i \geq 0$, and outside these finitely many charts g_B stays bounded above zero by a positive constant. Covering the compact family of supporting pairs (q, B) over Q by finitely many such neighbourhoods gives uniform constants c, C, A, a, γ and a uniformly bounded number of charts. Around *each* continued minimum z_i, place one more witness patch of radius t (even when $a_i > 0$, as that minimum may become a contact at a nearby query).

Now fix $q \in Q$ and a supporting ball B at q, and suppose all witness patches around its minimal contacts and continued minima z_i contain sites. Let a Delaunay simplex meet $B(q, \gamma t/2)$ at x_0, so $B(x_0, \gamma t/2) \subset B(q, \gamma t)$. Since the enclosing witness sites have power $g_B \leq At^2$, interpolating their heights on $B(q, \gamma t)$ gives $L_n \leq \ell_B + At^2$ on $B(x_0, \gamma t/2)$ (as in Fig. 4(a)). Lemma 3.3(b) then yields $g_B(p) \leq g_B(x) + Ct|p - x|$ for every simplex vertex p and every site x. Taking x as an enclosing witness site gives $g_B(p) = O(t)$ by compactness of V, falling below the positive outside-chart threshold for small t. Hence p lies in one of the charts, say around z_i. Taking x inside the radius-t patch at z_i (where $g_B(x) \leq a_i + Ct^2$) and setting $D = |p - z_i|$ yields

$$a_i + cD^2 \leq g_B(p) \leq a_i + C(t^2 + tD).$$

Crucially, a_i cancels from both sides, forcing $cD^2 \leq C(t^2 + tD)$ and hence $D \leq C't$ uniformly even as $a_i \to 0$. All vertices of any simplex meeting $B(q, \gamma t/2)$ thus fall inside the balls of radius $C't$ around the finitely many points z_i (the two arcs in Fig. 4(b)), with total sampling mass $O(t^\delta)$.

Covering the segment. For fixed $q \in Q$, putting $t = v^{1/\delta}$ for $1/n \leq v \leq v_0$ gives witness patches of mass $\geq av$ and a collector union of mass $\leq Av$ trapping every closed simplex meeting $D_n = B(q, \gamma n^{-1/\delta}/2)$. Lemma 3.1 bounds the expected count of such simplices by a uniform constant depending only on V, Q and (1.1). Any segment $s \subset Q$ is covered by at most $1 + C_Q|s|n^{1/\delta}$ balls of radius $\gamma n^{-1/\delta}/2$ along s, and summing their expected counts proves the first claim of Theorem 3.2.

Why an entire segment avoids the exceptional set. Suppose $E = (\operatorname{int} K) \setminus \mathcal{R}$ has zero $(d-1)$-dimensional Hausdorff measure. Fix a nonzero displacement vector $v \in \mathbb{R}^d$, and consider starting points $a \in \mathbb{R}^d$ for which $[a, a+v] \cap E \neq \varnothing$. Any such a has its orthogonal projection onto $v^\perp$ inside the orthogonal projection of E onto $v^\perp$. Since orthogonal projection is 1-Lipschitz, it sends an $\mathcal{H}^{d-1}$-null set to a $(d-1)$-dimensional Lebesgue null set in $v^\perp$. The cylinder of lines parallel to v through that projected null set therefore has d-dimensional Lebesgue measure zero, so for every $v \neq 0$ the set of bad starts a is Lebesgue-null. Integrating over v by Fubini's theorem and changing variables via $(a, v) \mapsto (a, b) = (a, a+v)$ shows that $[a, b] \cap E = \varnothing$ for almost every $(a, b) \in (\operatorname{int} K)^2$ (the diagonal $a = b$ has measure zero). By convexity of K, every segment between two interior points lies inside $\operatorname{int} K$; when it avoids E, it is a compact subset of $\mathcal{R}$, so the first part of the theorem applies with $Q = s$. □

One supporting ball along a segment. Suppose one fixed ball B supports every point of a fixed segment $s \subset Q \subset \mathcal{R}$. At scale t, take the 2δ witness patches around the minimal contacts for each endpoint of s using that ball B, plus one witness patch at each continued minimum z_i of g_B. Because the convex hull of those occupied sites encloses both $B(a, \gamma t)$ and $B(b, \gamma t)$, by convexity it encloses the entire capsule $s + B(0, \gamma t)$. Any simplex meeting $D_n = s$ contains a point $x_0 \in s$ with $B(x_0, \gamma t) \subset s + B(0, \gamma t)$, so Lemma 3.3(b) confines all its vertices into the fixed $O(t)$-balls around z_i. One call to Lemma 3.1 with $D_n = s$ then bounds $\mathbb{E}N_n(s) \leq C_Q$ uniformly for all $n \geq n_Q$.

A.2 Which contacts can a segment encounter?

For a generic smooth embedding $V = F(X) \subset \mathbb{R}^d$, we now determine how large the non-regular set $(\operatorname{int} K) \setminus \mathcal{R}$ can be. As shown on page 35, any set of zero $(d-1)$-dimensional Hausdorff measure $\mathcal{H}^{d-1}$ is missed by $[a, b]$ for almost every pair of interior endpoints. By contrast, a hypersurface of positive $\mathcal{H}^{d-1}$-measure (though still of d-dimensional volume zero) can split an open ball and block an open family of segments. The table below summarizes the five exceptional query sets analysed in Subsections A.2–A.4. Every configuration except the rank-nondeficient hypersurface at $d = j(\delta+1)$ lies in countably many smooth images of dimension at most $d-2$ (or is $\mathcal{H}^{d-1}$-null), leaving only the codimension-one interfaces of Subsections A.3–A.5.

Configuration at the query	Size of its image in query space
A degenerate contact	dimension at most $d-2$
A deficient j-tuple, $j \geq 2$	dimension at most $d-1-\lvert d-j(\delta+1)\rvert$
A rank drop of the secant map in the equality case	$\mathcal{H}^{d-1}$-null
Two extra contacts at the same sphere	dimension at most $d-2$ in the equality case
An extra contact in the base hyperplane	dimension at most $d-2$ in the equality case

Write $V = F(X)$, where X is a compact smooth δ-manifold without boundary, and vary F over the open set of smooth embeddings of full affine span in the C^∞ topology. Throughout our transversality arguments, marked sites $x_1, \ldots, x_m \in X$ are distinct, and we work inside countably many coordinate charts on tuples of distinct sites. Here the 3-*jet* of a smooth function at a point collects its value and partial derivatives up to order three in local coordinates. A smooth map f is *transverse* to a smooth constraint submanifold S if, at every point y with $f(y) \in S$, the image of df_y and $T_{f(y)}S$ span the ambient tangent space. By the parametric transversality theorem [19, Chapter II, Lemma 4.6, p. 53], whenever a family of maps depends transversely on auxiliary perturbation parameters, almost every small perturbation yields a transverse map whose preimage $f^{-1}(S)$ is a smooth submanifold of codimension $\operatorname{codim}(S)$. Below, the *secant map* takes a contact tuple and convex weights to their weighted mean in $\mathbb{R}^d$, j counts the contacts in a minimal representation of q, and a tuple is deficient when its tangents and chords fail to span $\mathbb{R}^d$.

A flat minimum costs two dimensions. Consider first a sphere of centre $c \in \mathbb{R}^d$ and radius $R > 0$, with power $u(x) = |F(x) - c|^2 - R^2$ on X. (A radius-zero sphere meets V only at its centre $q \in V$, whose restricted power already has positive definite Hessian $2|v|^2$, checked in Subsection A.3.) For $R > 0$, tangency at m distinct sites $x_1, \ldots, x_m$ imposes $m(\delta+1)$ equations: $u(x_i) = 0$ and $du(x_i) = 0$. The variables $x_1, \ldots, x_m, c, R$ provide $m\delta + d + 1$ coordinates. For a generic embedding, the solution manifold of these tangencies thus has dimension $(m\delta + d + 1) - m(\delta+1) = d + 1 - m$ (and is empty when $d + 1 - m < 0$). In particular, a generic supporting sphere touches V in at most $d+1$ points.

To establish both this count and the hyperplane counts below rigorously across all charts, we follow three standard steps: varying jets locally, perturbing densely on compact blocks, and intersecting countably many open dense conditions. Geometrically, each tangency or degeneracy imposes equations on the position and first three derivatives of V at the marked sites. Small smooth bumps supported near those sites act like independent local knobs that tune those derivatives freely, breaking accidental coincidences until every defining system achieves full row rank. Specifically, our constraints depend on the 3-jet of $u(x)$ and, when an auxiliary hyperplane $\nu \cdot p - h = 0$ appears, on the 3-jet of $\nu \cdot F(x) - h$, with sphere parameters (c, R) and plane parameters (ν, h) treated as free variables.

Varying the jets. A small perturbation Y of F changes $u(x)$ to first order by $2(F(x)-c)\cdot Y(x)$ and shifts $\nu \cdot F(x) - h$ by $\nu \cdot Y(x)$. For a positive-radius sphere, $F(x) - c \neq 0$ at every contact.

Whenever both sphere and hyperplane constraints appear together below, the hyperplane passes through at least two distinct sphere contacts and so cannot be tangent to the sphere at any contact. Near each marked site, the two normal vectors $F(x)-c$ and ν are therefore linearly independent. The matrix with these row vectors thus admits a smooth right inverse near each site. Applying that right inverse to arbitrary Taylor polynomials of degree at most three and multiplying by smooth cutoff functions with disjoint supports produces compactly supported perturbations Y. These Y realize any first-order variations of the 3-jets independently across all marked sites.

Obtaining density. Fix an embedding F and a compact block of distinct-site coordinate tuples and geometric parameters where the required linear-independence rank conditions hold. Because full row rank is open and the block is compact, finitely many perturbations Y_i span every first-order jet variation at every configuration in the block. For the finite-dimensional family $F + \sum_i b_i Y_i$, the derivative of the jet evaluation map with respect to (b_i) is therefore surjective near the block for small b_i. By parametric transversality [19, Chapter II, Lemma 4.6, p. 53], arbitrarily small perturbations make the jet map transverse to the smooth constraint submanifold in jet space. (When testing a degenerate contact, we adjoin a unit tangent vector $v \in S^{\delta-1}$ in $\ker(\operatorname{Hess} u)$ as an extra variable rather than stratifying the Hessian by rank.)

Obtaining a residual set. Exhaust the coordinate charts of distinct-site tuples and geometric parameters by countably many compact blocks, and exhaust each constraint submanifold in jet space by countably many compact subsets in its charts. Fix one compact block, one compact subset of the constraint submanifold, and an integer $k \geq 1$. At every point of the block mapping into that compact subset where the required normal row matrix (built from $F - c$ and, when present, ν) has smallest singular value at least $1/k$, we impose transversality of the jet map to the constraint submanifold. This condition is open in the C^∞ topology: any sequence of embeddings violating transversality at points meeting those closed bounds would yield a convergent subsequence of points whose limit violates the condition for F. It is also dense: apply the perturbation step above to the compact subset of the block where F satisfies the weaker singular-value bound $1/(2k)$. By continuity of jets and derivatives, for small perturbations that subset contains every configuration meeting the stronger bound $1/k$ for the perturbed embedding. Intersecting these open dense conditions over all countable blocks, constraint subsets and integers $k \geq 1$ gives the desired residual class.

Now count the dimension when one of the m contacts is a degenerate minimum. Adding a unit vector v in the kernel of the Hessian adds $\delta - 1$ variables. The equations $D^2u(v, \cdot) = 0$ impose δ independent constraints on second derivatives, and $D^3u(v, v, v) = 0$ adds one more (along a coordinate curve in direction v, a local minimum with vanishing second derivative cannot have nonzero third derivative). The net change is $(\delta - 1) - (\delta + 1) = -2$, even if the kernel is larger. Including at most $m - 1$ barycentric weights to form q puts these queries inside countably many smooth images of dimension at most $(d + 1 - m) - 2 + (m - 1) = d - 2$. Allowing weights to vanish ensures we exclude degenerate contacts even when they carry weight zero in representing q.

A missing direction imposes a second tangency. Suppose next that a minimal representing tuple of $j \geq 2$ contacts is deficient, meaning $W \neq \mathbb{R}^d$. Choose a unit normal $\nu \perp W$. Then all j contacts sit in one affine hyperplane $\nu \cdot p = h$ containing all j tangent spaces $T_{p_i}V$. We bound the dimension of such queries in two complementary ways. First, ignoring the sphere and counting only the j points on V with $j - 1$ barycentric weights gives at most $j\delta + j - 1 = j(\delta + 1) - 1$ parameters (the classical secant count [29, Introduction and Theorem 1.1]).

Second, project the sphere centre c orthogonally onto $\nu^\perp$ to obtain a new centre c'. For every site p in $\nu \cdot p = h$, Pythagoras gives $|p - c'|^2 - |p - c|^2 = 2h(\nu \cdot c) - (\nu \cdot c)^2$, independent of

p. Hence all j contacts remain equidistant from c', and since their tangents lie orthogonal to ν, the sphere centred at c' stays tangent to V at all j points. Projecting c onto $\nu^\perp$ removes one centre parameter: varying $\nu \in S^{d-1}$ ($d-1$ coordinates), $c' \in \nu^\perp$ ($d-1$ coordinates), $h \in \mathbb{R}$, and the new radius gives $2d$ geometric parameters. Both the sphere and the hyperplane are tangent to V at each of the j contacts, imposing $2j(\delta+1)$ equations whose gradients $F-c'$ and ν are independent because the hyperplane cuts the sphere in $j \geq 2$ distinct points. Adding $j\delta$ site coordinates and $j-1$ weights, then subtracting these $2j(\delta+1)$ equations, yields dimension at most

$$(j\delta + 2d) + (j-1) - 2j(\delta+1) = 2d - j(\delta+1) - 1.$$

Combining both bounds yields

$$\min\{j(\delta+1) - 1,\ 2d - j(\delta+1) - 1\} = d - 1 - |d - j(\delta+1)|.$$

Intersecting the residual sets over the finitely many contact counts gives embeddings where all bounds hold together. For $j = 1$, the query lies in V of dimension δ; for $j \geq 2$, the combined bound applies. In both cases, dimension reaches $d-1$ only when $d = j(\delta+1)$, while every smaller-dimensional set is $\mathcal{H}^{d-1}$-null and missed by almost every segment $[a, b]$. When $\delta+1$ does not divide d, Theorem 3.2 thus already yields Theorem 1.1; when $d = j(\delta+1)$, we must also handle crossings of the remaining $(d-1)$-dimensional interface.

A.3 When the query lies on the sampled hypersurface

When $\delta = d-1$, the manifold V is itself a hypersurface in $\mathbb{R}^d$, so an interior segment inside $\operatorname{int} K$ can pierce V directly (the case $j = 1$ of $d = j(\delta+1)$). Before treating general interfaces, we establish two facts for hypersurfaces: away from V, every interior query already has spanning contacts, and right on V, every query matches the $j = 1$ base-contact picture of Subsection A.4.

First take any interior query $q \in \operatorname{int} K \setminus V$. Any supporting representation of q must use at least two distinct contacts $p, x \in V$ of the supporting ball $B(c, R)$. Because $|x-c|^2 = |p-c|^2 = R^2$, expanding $|(x-p)+(p-c)|^2 = |p-c|^2$ yields $2(c-p)\cdot(x-p) = |x-p|^2 > 0$. Notice that $c-p$ is normal to the hypersurface V at p, and hence orthogonal to the $(d-1)$-dimensional tangent space T_pV. Thus the chord $x-p$ has a nonzero normal component and combines with T_pV to span all of $\mathbb{R}^d$. Let $E_{\deg}$ be the set of queries in $\operatorname{int} K$ admitting a supporting ball with a degenerate contact (shown in Subsection A.2 to be $\mathcal{H}^{d-1}$-null for a residual class of smooth embeddings). Every query in $(\operatorname{int} K) \setminus (V \cup E_{\deg})$ is therefore regular, so Theorem 3.2 applies to every compact subset of $(\operatorname{int} K) \setminus (V \cup E_{\deg})$.

Now suppose $q = p \in V \cap \operatorname{int} K$. For any convex combination $p = \sum_i \lambda_i p_i$ of points on V, expanding $|p_i - p|^2$ gives mean squared norm $\sum_i \lambda_i |p_i|^2 = |p|^2 + \sum_i \lambda_i |p_i - p|^2$. That sum attains the lower hull height $|p|^2$ only when every site p_i of positive weight equals p. Every supporting ball at p therefore touches V at p, forcing its centre to lie along the normal line $p + t\nu$ ($t \in \mathbb{R}$), where ν is a chosen unit normal to T_pV. The ball of radius $|t|$ centred at $p + t\nu$ supports p precisely when its power satisfies

$$g_t(x) = |x-p|^2 - 2t\,\nu\cdot(x-p) \geq 0 \qquad (x \in V). \tag{A.1}$$

At $t = 0$, the restricted Hessian of $g_0(x) = |x-p|^2$ on V equals $2|v|^2$ on every tangent vector $v \in T_pV$, so the radius-zero ball always has an ordinary contact at p. Furthermore, no other contact can lie in the tangent hyperplane $p + T_pV$, since any $x \neq p$ with $\nu\cdot(x-p) = 0$ satisfies $g_t(x) = |x-p|^2 > 0$ for every t. Point $p \in V \cap \operatorname{int} K$ therefore realizes the $j = 1$ case of the interval construction in Subsection A.4. To keep only nondegenerate crossings, we discard points on V belonging to a sphere with a degenerate contact (already $\mathcal{H}^{d-1}$-null) and points belonging to a sphere with two additional contacts (a 3-contact incidence of dimension at most

$d+1-3=d-2$ before projecting to p). In particular, for planar curves ($d=2, \delta=1$), both discarded sets have zero 1-dimensional Hausdorff measure.

A.4 One interval of centres at either kind of crossing

When $d=j(\delta+1)$, a generic segment can cross the $(d-1)$-dimensional interface of deficient queries where the contact directions span only a hyperplane W. Why does the local cell count still stay bounded across this crossing? We first discard three residual $\mathcal{H}^{d-1}$-null exceptions: a rank drop of the secant map, spheres with two extra contacts, and spheres with an extra contact in the base hyperplane. Once those are removed, the same geometric picture holds uniformly for every $j \geq 1$. At an interface point q, admissible sphere centres fill a nondegenerate closed interval on the line normal to W, and its two endpoint spheres each acquire one extra contact on opposite sides of $q+W$, supplying a clean two-sided coordinate system around the interface.

Write $q=\sum_{i=1}^{j} \alpha_i p_i$ for a minimal representation of $q \in \operatorname{int} K$ by contacts of a supporting sphere, with strictly positive weights $\alpha_i > 0$ summing to one. When $j=1$, q is simply a manifold point on the hypersurface V treated in Subsection A.3; when $j \geq 2$, we first remove the three exceptional sets just described.

Three exceptional sets. Let $x=(x_1,\dots,x_j) \in X^j$ satisfy $F(x_i)=p_i$, and let $\alpha=(\alpha_1,\dots,\alpha_j)$ satisfy $\sum_{i=1}^{j} \alpha_i = 1$. The secant map $\Phi(x,\alpha)=\sum_{i=1}^{j} \alpha_i F(x_i)$ sends a domain of dimension $j\delta+j-1=d-1$ into $\mathbb{R}^d$, and the image of its differential at (x,α) equals the span W of the contact tangents and chords. Wherever that differential has rank below $d-1$, its $(d-1)$-dimensional Jacobian vanishes. Covering the domain by countably many coordinate charts and applying the area formula shows that the image of all rank-drop points is $\mathcal{H}^{d-1}$-null. Removing these points ensures W is a $(d-1)$-dimensional hyperplane (for $j=1$, $\dim W=\delta=d-1$ holds already because F is an immersion).

Next exclude two residual incidence sets involving extra contacts (beyond the degenerate minima already excluded in Subsection A.2). First, any sphere touching V at the j base contacts and at two extra contacts belongs to an $(j+2)$-tangency family of dimension at most $d+1-(j+2)=d-j-1$. Adding the $j-1$ base weights to locate q puts these queries into countably many smooth images of dimension at most $d-2$. Second, exclude queries having an extra contact inside the affine hyperplane $q+W$. For $j=1$, no such contact exists by (A.1). For $j \geq 2$, count the sphere centre and radius plus the unit normal and offset of the hyperplane, giving $2d+1$ geometric parameters alongside $(j+1)\delta$ site coordinates and $j-1$ base weights. Both the sphere and the hyperplane are tangent to V at each of the j base sites ($2(\delta+1)$ equations per site), while at the extra site the sphere is tangent to V and the hyperplane passes through the site ($\delta+2$ equations). At every marked site, the two ambient gradients $F-c$ and ν are independent because the hyperplane cuts the sphere in at least two points, so jet perturbations again achieve transversality as in Subsection A.2. The net dimension is therefore at most

$$(j+1)\delta+2d+1-2j(\delta+1)-(\delta+2)+j-1=d-2.$$

Together with degenerate contacts, all exclusions discard only an $\mathcal{H}^{d-1}$-null set, which almost every interior segment misses.

Why the centres form an interval. Now take any remaining interface query q after these exclusions, with j base contacts $p_1,\dots,p_j$ and positive weights $\alpha_1,\dots,\alpha_j$. Every affine support attaining $L_V(q)$ must touch V at every p_i, because all gaps are nonnegative and average to zero at q. Equal squared distances from the j base contacts and tangency along $T_{p_i}V$ impose $j-1+j\delta=d-1$ independent linear equations on the sphere centre c, whose row span equals

W. Letting ν be a unit normal to W, all solutions c therefore lie on one normal line. Pick a supporting centre c_0 at q and parametrize this line by $c(a) = c_0 + \frac{1}{2}a\nu$ ($a \in \mathbb{R}$), where the factor $1/2$ makes a the linear slope of the affine support in direction ν. The power of the sphere centred at $c(a)$ passing through $p_1, \ldots, p_j$ then reads

$$g_a(p) = |p|^2 - |p_1|^2 - 2c_0 \cdot (p - p_1) - a\,\nu \cdot (p - p_1).$$

By the compactness proof in Subsection A.1, the centres of all supporting balls at an interior query q remain bounded. Requiring $g_a(p) \geq 0$ on all of V therefore cuts out a nonempty compact interval $[a_-, a_+]$ on the normal line.

At either endpoint $a \in \{a_-, a_+\}$, the sphere must touch V at some contact outside the j base contacts. Otherwise, positive definite Hessians at the ordinary base contacts and a positive minimum on the compact remainder of V outside small neighbourhoods of the base contacts would keep $g_{a'} \geq 0$ slightly past the endpoint. By our second exclusion, each endpoint sphere has a unique extra contact. As a varies near that endpoint, the implicit function theorem continues the local minimum of g_a near the extra contact to a unique C^1 point $p_*(a)$ (smooth when V is C^∞) with derivative

$$\frac{d}{da} g_a(p_*(a)) = -\nu \cdot (p_*(a) - p_1) \neq 0,$$

because $\nabla_p g_a(p_*(a)) = 0$ along V and our third exclusion keeps $p_*(a) \notin q + W = p_1 + W$. Moving a toward the side where $g_a(p_*(a)) > 0$ produces valid supporting balls, proving $a_- < a_+$. Moreover, $-\nu \cdot (p_*(a) - p_1)$ must be positive at a_- and negative at a_+, placing the two endpoint extra contacts $p_+ = p_*(a_+)$ and $p_- = p_*(a_-)$ on strictly opposite sides of $q + W$ (matching the two diamonds in Fig. 6(b)). Finally, no extra contact can occur at any interior parameter $a \in (a_-, a_+)$. Because $g_a(p)$ is affine in a and nonnegative on $[a_-, a_+]$, vanishing at an interior point would force it to vanish for all a, placing p inside $q + W$. Any interior supporting ball at q thus forces every convex representation of q attaining the lower height $L_V(q)$ to use only the affinely independent base contacts with their unique weights $\alpha_1, \ldots, \alpha_j$.

How the interval serves nearby queries. Because the $d - 1$ linear equations for the sphere centre at (x, α) have full rank $d - 1$, a reference centre $c_0(x)$ and unit normal $\nu(x)$ vary in a C^1 manner (smoothly when V is C^∞) for base tuples x' near x. Any fixed interior parameter $a \in (a_-, a_+)$ stays strictly supporting for nearby tuples x' by positive definiteness at the base contacts and compactness away from the base charts. At each endpoint $a_\pm$, the nonzero derivative $\frac{d}{da} g_a(p_*(a)) \neq 0$ lets the implicit function theorem continue both $a_\pm(x')$ and $p_\pm(x')$ as C^1 functions of x' (smooth when V is C^∞), keeping both endpoint spheres empty and preventing extension past $a_\pm(x')$.

On a small coordinate chart around (x, α), the secant map Φ is a C^1 embedding (smooth when V is C^∞) onto the $(d-1)$-dimensional interface (simply V itself when $j = 1$) with tangent hyperplane W. Adjoining either extra contact $p_+(x')$ or $p_-(x')$ with weight $\mu \in \mathbb{R}$ gives the two maps

$$(x', \alpha', \mu) \longmapsto (1 - \mu)\Phi(x', \alpha') + \mu p_\pm(x').$$

At $\mu = 0$, the μ-derivative is $p_\pm(x') - \Phi(x', \alpha')$, which points strictly off W in opposite directions for $+$ and $-$. Both maps are therefore locally invertible near $(x, \alpha, 0)$, and for small $\mu > 0$ they parametrize the two open sides adjacent to the interface. Each point $y = (1 - \mu)q' + \mu p_\pm(x')$ with small $\mu > 0$ is supported by the endpoint sphere at $a_\pm(x')$. In fact, that ball is the *unique* supporting ball at y, because positive base weights confine any supporting centre to the normal line at x' and vanishing gap at $p_\pm(x')$ selects $a_\pm(x')$. Since the j base contacts and $p_\pm(x')$ are affinely independent, every convex representation of y by contacts of its supporting ball (equivalently, attaining $L_V(y)$) has the same unique weights. Shrinking the chart bounds the

derivatives of both inverse maps and all contact data on each side while sharing the base tuple at $\mu = 0$; in particular, μ is comparable to the Euclidean distance from y to the interface, giving the two-sided coordinates used in Subsection A.5.

A.5 Counting cells across a contact change

As a query y approaches an interface of contact change from either side, the weight μ on the extra contact p_+ or p_- shrinks to zero. A ball around y of radius $bn^{-1/\delta}$ can therefore straddle the interface and touch both sides at once. How do we confine cell vertices uniformly on both sides using a single ball radius? Proposition A.1 pairs the two extreme supporting spheres at the interface into a convex crease $L_\vee = \max(\ell_{a_-}, \ell_{a_+})$ (Fig. 5(b)). Placing witness patches around the base contacts and both extra contacts builds two cones meeting along a disk in $q + W$ (Fig. 13). Those two cones trap L_n within $O(t^2)$ of $L_\vee$ on an entire ball around y, letting Lemma 3.3(c) pull every cell vertex back into the $j + 2$ contact charts.

Proposition A.1. *Fix integers $\delta, j \geq 1$ and set $d = j(\delta + 1)$. Let $V \subset \mathbb{R}^d$ be a compact C^2 embedded δ-manifold without boundary and of full affine span, and put $K = \operatorname{conv} V$. Draw n sites independently on V from a common law satisfying* (1.1).

Suppose $q_0 \in \operatorname{int} K$ has a minimal representation by j contacts of a supporting ball, and every contact of every supporting ball at q_0 is ordinary. Assume further that the chords and tangent spaces of the representing contacts span a hyperplane W, that the centres supporting q_0 form a nontrivial closed interval, and that each endpoint ball has exactly one extra contact, with these two contacts lying on opposite sides of $q_0 + W$.

Then there are an open neighbourhood $U \subset \operatorname{int} K$ of q_0, constants $b, C > 0$ and $n_0 \in \mathbb{N}$, depending only on V, q_0 and (1.1)*, such that for every $y \in U$ fixed before sampling and every $n \geq n_0$, the expected number of closed d-dimensional Delaunay simplices meeting $B(y, bn^{-1/\delta})$ is at most C.*

Proof. Observe first that these hypotheses match Subsection A.4. Indeed, no interior parameter along the interval of centres at q_0 can gain an extra contact. Because its power is affine and nonnegative on the interval, vanishing at an interior point would make it vanish at both endpoints and force the same unique extra contact on both sides of $q_0 + W$. Thus the local two-sided coordinate system of Subsection A.4 applies around q_0 (with C^1 regularity under our C^2 assumption on V). On the interface, both extreme affine supports agree at all base contacts, forming the V-shaped crease drawn in Fig. 5(b).

Fix a query $y \in U$ near q_0, writing $y = (1-\mu)q + \mu p_\star$ with q on the interface, $p_\star \in \{p_+, p_-\}$ the extra contact on the side of y (taking either if $\mu = 0$), and $\mu \geq 0$ small. On a compact chart around q_0, the j base contacts $p_1, \dots, p_j$ at q have weights bounded uniformly away from zero. Translate the origin of displacements to q and pick a unit normal ν to $q + W$. Write $z(x) = \nu \cdot (x - q)$ for signed height above $q + W$, labelling p_+ on the positive side and p_- on the negative side. Shrinking U ensures $z(p_+) \geq b_0 > 0$ and $z(p_-) \leq -b_0 < 0$ uniformly. Let ℓ_0 correspond to the midpoint of the interval of centres at q, so supporting planes through the base contacts read $\ell_a = \ell_0 + az$ for $a \in [a_-, a_+]$, with $\ell_{a_\pm}$ touching $p_\pm$. Define the gap $g_a(x) = |x|^2 - \ell_a(x)$ and the two-plane crease $L_\vee = \max(\ell_{a_-}, \ell_{a_+})$, which equals ℓ_{a_+} on $\{z \geq 0\}$ and ℓ_{a_-} on $\{z \leq 0\}$.

A sampled disk in the interface plane. At scale $t > 0$, place the 2δ signed witness patches of radius proportional to t around each base contact $p_1, \dots, p_j$ (as in Subsection A.1), plus one patch of radius t around p_+ and around p_-. By (1.1), each patch has sampling mass at least ct^δ. Suppose every patch contains a site, and choose one representative X_i from each base patch and $X_\pm$ from the extra patches. By tangency of both endpoint spheres at p_i, every base representative satisfies $z(X_i) = O(t^2)$ and $g_{a_\pm}(X_i) = O(t^2)$. Applying the

enclosing argument of Subsection A.1 inside the $(d-1)$-dimensional hyperplane W shows that the orthogonal projections πX_i onto $q+W$ enclose a $(d-1)$-dimensional disk of radius $c_1 t$ centred at q in $q+W$. Meanwhile, $z(X_+) \geq b_0/2$ and $z(X_-) \leq -b_0/2$, with $g_{a_+}(X_+) = O(t^2)$ and $g_{a_-}(X_-) = O(t^2)$.

Next build points inside $q+W$ where L_n stays bounded above by $L_\vee + Ct^2$. Join each base representative X_i having $z(X_i) \neq 0$ by a segment to the extra representative X_- or X_+ of opposite sign (drawn for $z(X_i) > 0$ in Fig. 13(a)), and mark where it crosses $q+W$. Because $|z(X_i)| = O(t^2)$ while $|z(X_\mp)| \geq b_0/2$, the crossing point puts weight $O(t^2)$ on $X_\mp$. It therefore stays within $O(t^2)$ of πX_i, and its interpolated sample height exceeds both ℓ_{a_+} and ℓ_{a_-} by at most Ct^2 (the bounded gap at $X_\mp$ is multiplied by $O(t^2)$). When $z(X_i) = 0$, X_i already lies in $q+W$ with gap $O(t^2)$. Replacing πX_i by points within $O(t^2)$ preserves enclosure for small t, so convex combinations of these cuts produce a $(d-1)$-dimensional disk $D \subset q+W$ of radius $c_1 t/2$ centred at q where $L_n \leq L_\vee + Ct^2$.

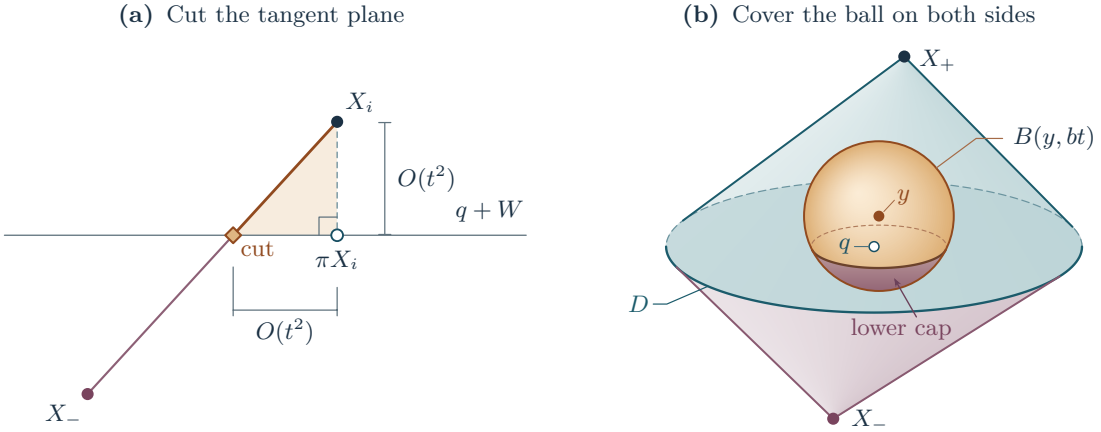


Figure 13: Enclosing a two-sided ball around y. In (a), a base site X_i above the interface tangent plane $q+W$ is connected to the opposite extra site X_-, cutting $q+W$ at the diamond within $O(t^2)$ of the orthogonal projection πX_i (offsets enlarged). In (b), the outer ellipse shows the common base disk D seen obliquely and the inner ellipse shows the section of $B(y, bt)$ in $q+W$; when $B(y, bt)$ straddles $q+W$ from above, the upper cone with tip X_+ covers the upper part while the lower cone with tip X_- covers the shaded cap below $q+W$.

Extending the disk to both sides. Form the two cones $\mathrm{conv}(D \cup \{X_+\})$ and $\mathrm{conv}(D \cup \{X_-\})$, which lie on opposite sides of the interface hyperplane $q+W$. On $\{z \geq 0\}$ we have $L_\vee = \ell_{a_+}$ and on $\{z \leq 0\}$ we have $L_\vee = \ell_{a_-}$, while L_n is convex and satisfies $L_n(X_\pm) \leq \ell_{a_\pm}(X_\pm) + Ct^2$ and $L_n \leq L_\vee + Ct^2$ on D. Because $L_\vee$ is affine on each half-space separately, convexity inside each cone yields

$$L_\vee(x) \leq L_n(x) \leq L_\vee(x) + Ct^2 \tag{A.2}$$

throughout the union $\mathrm{conv}(D \cup \{X_+\}) \cup \mathrm{conv}(D \cup \{X_-\})$.

To show that this union of two cones contains $B(y, bt)$ for a uniform $b > 0$ once U (and hence μ) is small, set $\widetilde{y} = (1-\mu)q + \mu X_\star$, so $|\widetilde{y} - y| = \mu |X_\star - p_\star| \leq C\mu t$. Radial projection from the tip $X_\star$ onto $q+W$ sends $\widetilde{y}$ directly to q. Because $|z(X_\star)| \geq b_0/2$ while $|z(y)| = \mu |z(p_\star)|$ is small, this radial projection has uniformly bounded derivative near y. Any point of $B(y, bt)$ on the $X_\star$-side of $q+W$ thus projects into $q+W$ within distance $C(b+\mu)t$ of q, landing inside D for small enough μ and b while keeping height strictly below $|z(X_\star)|$. If $B(y, bt)$ crosses $q+W$ into the opposite shaded cap (Fig. 13(b)), then $\mu |z(p_\star)| = |z(y)| \leq bt$, forcing $|y - q| \leq Cbt$;

projecting points of that cap from the opposite tip likewise lands in D. Shrinking U first to bound μ and then fixing $b > 0$ yields (A.2) on all of $B(y, bt)$, even across the interface.

Where the vertices can lie. Suppose a closed Delaunay simplex meets $B(y, bt/2)$ at x_0, so $B(x_0, bt/2) \subset B(y, bt)$ where (A.2) holds. Lemma 3.3(c) yields a parameter $a \in [a_-, a_+]$ such that every simplex vertex p and every sampled site x obey

$$g_a(p) \le g_a(x) + Ct|p - x|.$$

Choosing $x = X_i$ as any base representative (with gap $O(t^2)$ for both endpoint supports and thus for ℓ_a) gives $g_a(p) = O(t)$ by boundedness of V. By compactness of $[a_-, a_+]$ and of the query chart around q_0, g_a has a uniform positive lower bound outside the $j + 2$ contact charts around $p_1, \dots, p_j, p_+, p_-$. For small t, every vertex p therefore enters one of those charts.

If p falls in a base chart around p_i, ordinary tangency gives $g_a(p) \ge c|p - p_i|^2$ uniformly over $a \in [a_-, a_+]$. Comparing against the base site within Ct of p_i yields $c|p - p_i|^2 \le C(t^2 + t|p - p_i|)$ and hence $|p - p_i| = O(t)$, just as in Subsection A.1. If instead p falls in the chart around p_+, where $z \ge b_0/2$, rewrite $g_a = g_{a_+} + (a_+ - a)z$. Since $g_{a_+} \ge 0$, having $g_a(p) = O(t)$ and $z(p) \ge b_0/2$ forces $a_+ - a = O(t)$. When comparing $g_a(p)$ with $g_a(X_+)$, switching to g_{a_+} alters the difference by $(a_+ - a)|z(p) - z(X_+)| \le Ct|p - X_+|$. Ordinary tangency of g_{a_+} at p_+ then gives $c|p - p_+|^2 \le C(t^2 + t|p - p_+|)$ and $|p - p_+| = O(t)$, and reversing the sign of z handles p_- identically. Every vertex of a simplex meeting $B(y, bt/2)$ thus lies in a deterministic union of $j + 2$ collector balls of radius $O(t)$ and mass $O(t^\delta)$. Putting $t = v^{1/\delta} \ge n^{-1/\delta}$ satisfies Lemma 3.1 for $B(y, bn^{-1/\delta}/2)$, and halving b completes the proof. □

Completion of the proof of Theorem 1.1. By Subsection A.1, the regular queries form an open set $\mathcal{R} \subset \operatorname{int} K$. For each query point q_0 on a nondegenerate crossing interface of Subsection A.4, Proposition A.1 yields an open neighbourhood $U \subset \operatorname{int} K$ where the local expected crossing bound holds uniformly at scale $bn^{-1/\delta}$. Let $G \subset \operatorname{int} K$ be the union of $\mathcal{R}$ and all such neighbourhoods U, so G is open. By the dimension bounds and exclusions in Subsections A.2–A.4, the exceptional complement $E = (\operatorname{int} K) \setminus G$ has zero $(d-1)$-dimensional Hausdorff measure. Applying the projection argument on page 35 directly to E shows that for Lebesgue-almost every pair $(a, b) \in (\operatorname{int} K)^2$, the segment $[a, b]$ avoids E and therefore stays inside G (possibly crossing interfaces of contact switch there).

To obtain uniform constants on any compact subset $Q \subset G$, cover Q by finitely many open neighbourhoods on which the local ball bound of Subsection A.1 or Proposition A.1 holds. Choosing the smallest radius coefficient b, the largest local expectation constant C, and the largest threshold n_0 from this finite cover gives uniform parameters across Q. Any fixed segment $s \subset Q$ is covered by at most $O_Q(1 + |s|n^{1/\delta})$ balls of radius $bn^{-1/\delta}$ centred along s, and summing their expected cell counts establishes the bound on Q. For any fixed segment $s = [a, b]$ avoiding E, taking $Q = s$ gives the individual-segment bound as well. □

Notice that the final finite-cover argument also yields a deterministic criterion for any specific compact smooth embedding $V \subset \mathbb{R}^d$ of full affine span. Suppose one checks directly that, outside a set of zero $(d-1)$-dimensional Hausdorff measure, every interior query either is regular or satisfies the hypotheses of Proposition A.1. Then the expected crossing bound holds uniformly on every compact subset of the resulting open set G.

A.6 Contact points in the surface and curve experiments

We now verify the contact geometry of the two experimental families in Section 6. First, we show that the test segments traced on the quadratic and quartic embeddings in $\mathbb{R}^4$ and $\mathbb{R}^5$ lie

inside the regular set $\mathcal{R}$ of Theorem 3.2. Next, we explain why the moving sphere touches the sampled smooth curve Γ_∞ at three points and the analytic comparison curve Γ_5 at four points across the common experimental window $|\tau| \le 1/32$ in Fig. 11.

Moving ordinary contacts in $\mathbb{R}^4$ and $\mathbb{R}^5$. For the six surface and submanifold series in Fig. 10(a), take the δ-dimensional parameter manifold $x^2 + |u|^2 + \kappa|u|^4/2 = 1$ in $\mathbb{R} \times \mathbb{R}^\delta$, where $\kappa = 0$ gives the quadratic profile ($\delta \in \{2, 3\}$) and $\kappa = 1$ gives the quartic profile ($\delta = 2$). For each ambient dimension $d \in \{4, 5\}$, put $m = d - \delta - 1 \in \{0, 1, 2\}$ and fix the unit vector $w = (2, 1, 0, \dots, 0)/\sqrt{5} \in \mathbb{R}^\delta$. Define the orthonormal-row linear projection $P\colon \mathbb{R}^\delta \to \mathbb{R}^m$ by omitting the block for $m = 0$, taking $P = w^\top$ for $m = 1$, and taking $P = I_2$ for $m = \delta = 2$. We embed the manifold into $\mathbb{R}^d$ via

$$H(x, u) = (x, 2u, xPu),$$

and sample sites over the whole image by rejection using the exact δ-dimensional volume Jacobian of H. Along the query segment parameterized by $|t| \le \sqrt{5}/16$, write $v = tw$ and $a = \sqrt{1 - |v|^2 - \kappa|v|^4/2}$. The query point and its two candidate contacts on the embedded manifold are then

$$q(t) = (0, 2v, 0), \qquad p_+(t) = H(a, v), \qquad p_-(t) = H(-a, v).$$

Notice that $p_+(t)$ and $p_-(t)$ share the same middle block $2v$, while their first and last blocks cancel at their midpoint $q(t)$. At $t = 0$, these contacts are the images of $(\pm 1, 0)$, called the poles; at every nonzero t along the segment, $u = v \neq 0$ stays away from the poles.

To prove that $p_\pm(t)$ are the unique contacts of a supporting ball at $q(t)$, put $\sigma = |v|^2 \le 5/256$ and $\rho = |u|^2 \le 1$. When all coordinates of xu are kept ($P = I_\delta$), using $x^2 = 1 - \rho - \kappa\rho^2/2$ simplifies the squared norm of $H(x, u)$ to a polynomial $f(\rho) = |H(x, u)|^2 = 1 + 4\rho - (1 + \kappa/2)\rho^2 - \kappa\rho^3/2$. For the query $q(t)$, take the sphere through $p_+(t)$ and $p_-(t)$ with centre $c(t) = (0, f'(\sigma)v/2, 0)$. Since $H(x, u) \cdot c(t) = f'(\sigma)(u \cdot v)$, its power at any point $H(x, u)$ of the full embedding equals

$$|H(x,u)|^2 - |H(a,v)|^2 - 2\big(H(x,u) - H(a,v)\big)\cdot c(t) = f'(\sigma)|u-v|^2 - \big[1+\kappa(1+\rho+2\sigma)/2\big](\rho-\sigma)^2.$$

When P projects onto a proper subspace whose row space contains w (including $m = 1$), dropping the orthogonal complement lowers the squared norm at $H(x, u)$ by $x^2|(I - P^\top P)u|^2$. Because w lies in the row space of P, we have $(I - P^\top P)v = 0$ and thus $x^2|(I - P^\top P)u|^2 = x^2|(I - P^\top P)(u - v)|^2 \le |u - v|^2$. Moreover, $|\rho - \sigma| = ||u|^2 - |v|^2| \le (|u| + |v|)|u - v| \le (1 + \sqrt{\sigma})|u - v|$. Over $\sigma \le 5/256$ and $\rho \le 1$, we have $f'(\sigma) = 4 - (2 + \kappa)\sigma - \frac{3}{2}\kappa\sigma^2 > 3.94$, the bracket satisfies $1 + \kappa(1 + \rho + 2\sigma)/2 \le 2.02$, and $(1 + \sqrt{\sigma})^2 < 1.30$. Even after removing the dropped coordinates, the power therefore exceeds $(3.94 - 1 - 2.02 \cdot 1.30)|u - v|^2 > 0.3|u - v|^2$ for all $u \neq v$. When $m = 0$ (with $\kappa = 0$), $|H(x, u)|^2 = 1 + 3\rho$, so the sphere centred at $(0, 3v/2)$ has exact power $3|u - v|^2$. In every case, this quadratic lower bound isolates $p_+(t)$ and $p_-(t)$ as the sole contacts of the ball and proves positive definite Hessians in the local chart coordinate u on each sheet $x = \pm\sqrt{1 - |u|^2 - \kappa|u|^4/2}$.

It remains to verify that the contact directions span $\mathbb{R}^d$ and that no other supporting ball exists at $q(t)$. Differentiating $H(\pm\sqrt{1 - |u|^2 - \kappa|u|^4/2}, u)$ along any velocity $h \in \mathbb{R}^\delta$ at $u = v$ and averaging the two sheets gives $(0, 2h, 0)$. Taking half their difference modulo the chord $p_+(t) - p_-(t) = (2a, 0, 2aPv)$ cancels the first coordinate and leaves $(0, 0, aPh)$. Since P has full rank m and $a > 0$, these tangent vectors and the chord span all $1 + \delta + m = d$ ambient directions, ruling out any supporting hyperplane to $\operatorname{conv} V$ through $q(t)$ as well. Finally, any other supporting ball at $q(t)$ has nonnegative gaps at $p_+(t)$ and $p_-(t)$ averaging to zero at their midpoint $q(t)$, so it must touch both contacts. By the spanning property, tangency and equidistance at those two points pin down the centre uniquely. Every point along the tested segment is therefore regular, and Theorem 3.2 applies directly.

The three contacts of Γ_∞ and four contacts of Γ_5. To interpret Fig. 11, we compare the exact contact sets of the two sampled curves Γ_∞ and Γ_5 with the moving sphere of radius 2 centred at $H_\tau(-1,0,0)$ across the common experimental phase window $|\tau| \le 1/32$. Both curves were tested along the vertical segment with endpoints $q_* \pm e_z/32$, where $e_z = (0,0,1)$ and $q_* = (-1/3, \sqrt{2}/6, \sqrt{2}/6)$ lies inside the contact triangle BR_+R_- (whose rotated contact planes sweep phases $|\tau| < 1/40$ across the segment). Recall the screw motion H_u that rotates the first two coordinates by angle u and translates the third by u. The orbit of $B = (1,0,0)$ is the front helix of Subsection C.1, whose osculating circle at B is $P_R(\theta) = (-1+R\cos\theta, R\sin\theta/\sqrt{2}, R\sin\theta/\sqrt{2})$ at $R = 2$. Both curves take the shared form $\gamma(v) = H_{h(v)}P_{R(\theta(v))}(\theta(v))$ for $v \in \mathbb{R}/(2\pi\mathbb{Z})$ with $h(v) = \sin(2v)/4$, and both were sampled uniformly by arclength via acceptance-rejection on parameter proposals.

For the explicitly sampled smooth curve Γ_∞ realizing Corollary 5.2 (see `contact_curves/geometry.md` for the parametrizations), smooth flat step transitions hold $\theta(v)$ constant at $0, \pi/2, \pi, 3\pi/2$ on four plateaus $|v - j\pi/2| \le \pi/12$, across each of which $h(v)$ spans $[-1/8, 1/8]$. Setting $R = 2$ for $\cos\theta \ge 0$ leaves the three helical arcs H_hB, H_hR_+, H_hR_- exact on their plateaus, while setting $R(\theta) = 2 + \exp(1/\cos\theta)/4$ for $\cos\theta < 0$ lifts the fourth plateau (at $\theta = \pi$) outward by $R - 2 = 1/(4e)$. Across the common window $|\tau| \le 1/32$, this lift raises the sphere power strictly above zero on the fourth plateau (by at least $3.8/(4e)$), while the connecting transitions have $|h| \ge 1/8$ (hence $|h - \tau| \ge 3/32$) and keep power bounded below by $0.018(3/32)^4 > 0$. Consequently, the moving sphere touches Γ_∞ at exactly three points for every $|\tau| \le 1/32$.

For the analytic comparison curve Γ_5, we instead keep $R \equiv 2$ and set

$$u = \sin(2v)/4, \quad \theta(v) = v - \tfrac{1}{3}\sin(4v) + \tfrac{1}{24}\sin(8v), \quad \Gamma_5(v) = H_uP_2(\theta(v)) \quad (0 \le v < 2\pi).$$

On Γ_5, the front arc and two rear arcs correspond to intervals of radius $\pi/12$ centred at $v = 0, \pi/2, 3\pi/2$ modulo 2π, across each of which u spans $[-1/8, 1/8]$. Near $v = 0$, differentiating gives $\theta'(v) = (8/3)\sin^4(2v)$, keeping $|\theta(v)| < 8.3 \times 10^{-6}$ whenever $|u| \le 1/32$.

We now verify that for every phase $|\tau| < 1/4$ (and in particular throughout $|\tau| \le 1/32$), the moving sphere of radius 2 centred at $H_\tau(-1,0,0)$ touches Γ_5 at four points. Apply the inverse screw isometry $H_{-\tau}$ and put $w = u - \tau$, $a = 1 - \cos w$, and $b = w - \sin w$. The sphere returns to centre $(-1,0,0)$ and radius 2, and its power at a swept point $H_uP_2(\theta)$ equals

$$w^2 + 2a - 4a\cos\theta + 2\sqrt{2}\,b\sin\theta.$$

Minimizing over $\theta \in [0, 2\pi)$ gives $w^2 + 2a - \sqrt{16a^2 + 8b^2}$, which vanishes at $w = 0$. For $w \ne 0$ it stays strictly positive: the difference of squares $D(w) = (w^2 + 2a)^2 - 16a^2 - 8b^2$ is even, satisfies $D(0) = 0$, and has derivative $D'(w) = 4(w + \sin w)(w^2 - 2a) > 0$ for $w > 0$ since $2a = 4\sin^2(w/2) < w^2$. Thus the sphere meets the swept surface precisely along the circle $H_\tau P_2$ at $u = \tau$. Whenever $|\tau| < 1/4$, the equation $\sin(2v) = 4\tau$ has four distinct solutions v modulo 2π. Because $\theta(v)$ increases strictly over a period, they map to four distinct contacts on $H_\tau P_2$; at $\tau = 0$, alongside B, R_+, R_-, the fourth contact at $v = \pi$ is $(-3,0,0)$.

B Why a contact slice confines a cell

To establish Proposition 4.1 and Corollary 4.2, we formalize the axial localization pictured in Fig. 7(b). Lemma B.1 places witness patches on both sides of any interior axial level v so that, once occupied, every simplex reaching a thin axial window around $q = (v, y_q)$ takes all its vertices from a narrow belt $V \cap \{|x - v| \le Dt\}$. Counting sites in one belt gives the upper bound via Lemma 3.1, while spacing marks $3Dt$ apart in axial coordinate separates their belts completely and yields the linear lower bound. We conclude by checking the interior-simplex property and constructing the higher-dimensional spherical factors for Corollary 4.2.

Lemma B.1. *Let $d \geq 2$ and $1 \leq \delta \leq d-1$ be integers, and adopt the δ-manifold $V \subset \mathbb{R}^d$, axial height h, compact set Q, open set I, and sampling law of Proposition 4.1, with its hypotheses (1) and (2). All patches and belts below stay inside $V \cap x^{-1}(I)$. Then there exist positive constants ρ, D, t_0, c, C and an integer J with the following property. For every $q = (v, y_q) \in Q$ and $0 < t \leq t_0$, define the rectangular window*

$$U_t(q) = \{(x, y) : |x - v| \leq \rho t,\ |y - y_q| \leq \rho\}.$$

We can choose at most J deterministic patches on V, each of sampling probability at least ct, with the following property. Whenever every patch contains a site, $U_t(q)$ lies in the interior of the sample hull, and every vertex of every closed full-dimensional projected lower cell meeting $U_t(q)$ lies on the belt $V \cap \{|x - v| \leq Dt\}$. Moreover, this belt has sampling probability at most Ct.

Proof. Enclosing the window. Write points in $\mathbb{R}^d$ as $p = (x, y)$ with $y \in \mathbb{R}^{d-1}$, and fix $q = (v, y_q) \in Q$. By hypothesis (2) on the slice $V_v = \{p \in V : x(p) = v\}$, y_q lies inside the convex hull of the y-coordinates of V_v. We can therefore pick finitely many slice points $(v, y_j) \in V_v$ whose y-coordinates enclose a ball around y_q (for instance, expressing the vertices of a small simplex around y_q as convex combinations from V_v).

Because x has no critical points on $V \cap x^{-1}(I)$ (its differential never vanishes along V), the submersion theorem lets us use x itself as the first local coordinate in charts on V, with a bounded positive Jacobian converting coordinate volume to intrinsic volume on shrunk charts. Covering the compact preimage of a closed neighbourhood of $x(Q)$ in I by finitely many charts bounds the intrinsic volume of each belt $V \cap \{|x - v| \leq t\}$ by $O(t)$, giving sampling probability $O(t)$ by the upper density bound. Similarly, inside a chart around (v, y_j), taking a small transverse ball together with either axial interval $|x - (v \pm t)| < \eta t$ creates a witness patch of intrinsic volume at least a fixed multiple of t and sampling probability at least $c_\eta t$. Once the transverse ball and t are small enough, every patch point satisfies $|y - y_j| < \eta$.

In the rescaled coordinates $((x - v)/t, y)$, the ideal points $(\pm 1, y_j)$ enclose the product $[-1, 1] \times \operatorname{conv}\{y_j\}$. Choosing one site from each witness patch moves these points by at most 2η in the rescaled coordinates, shifting the support function in any unit direction by at most 2η. Fixing $\rho, \eta > 0$ small enough ensures that the convex hull of the picked sites still contains the doubled window $\{(x, y) : |x - v| \leq 2\rho t,\ |y - y_q| \leq 2\rho\}$. Because the charts continue (v, y_j) smoothly to nearby levels, covering compact Q by finitely many neighbourhoods makes J, c, ρ, t_0 uniform.

Keeping the vertices on the belt. Consider the affine support $\ell(p) = h(v) + h'(v)(x(p) - v)$ defined by the tangent line to h at v. By convexity of h on $[x_-, x_+]$, ℓ supports all lifted heights, and Taylor's formula bounds its gap $g(p) = h(x(p)) - \ell(p)$ between $c_0(x(p) - v)^2$ and $C_0(x(p) - v)^2$ with $c_0 = m/2$ and $C_0 = M/2$. On the doubled window, interpolating the occupied witness sites (each having gap at most $C_1 t^2$) bounds $L_n \leq \ell + C_1 t^2$.

Here the elongated window shape is crucial: a cell vertex p might sit far across the body at Euclidean distance of order one, where applying Lemma 3.3(b) to an isotropic ball of radius ρt inside $U_t(q)$ would give only $|x(p) - v| = O(\sqrt{t})$. Instead, the wide transverse radius ρ lets us step toward p using Lemma 3.3(a). Suppose a closed full-dimensional lower cell meets $U_t(q)$ at q' and has a vertex p with $|x(p) - v| = Lt$ and $L \geq 1$. Moving a fraction α/L of the way from q' toward p changes the axial coordinate by at most $\alpha(1 + \rho)t$ and the transverse coordinates by at most $\alpha \operatorname{diam}(\operatorname{conv} V)$. Fixing $\alpha \in (0, 1)$ small keeps this point inside the doubled window (axial half-width $2\rho t$, transverse half-width 2ρ). Applying Lemma 3.3(a) with $\lambda = \alpha/L$ gives

$$c_0 L^2 t^2 \ \leq\ g(p) \ \leq\ \frac{L}{\alpha} C_1 t^2.$$

Because the left side grows quadratically in L while the right side grows linearly, L cannot exceed $C_1/(c_0\alpha)$. Putting $D \geq \max\{1, C_1/(c_0\alpha)\}$ confines every vertex into $V \cap \{|x-v| \leq Dt\}$, and shrinking t_0 keeps the belt inside the noncritical strip where its mass is $O(t)$. □

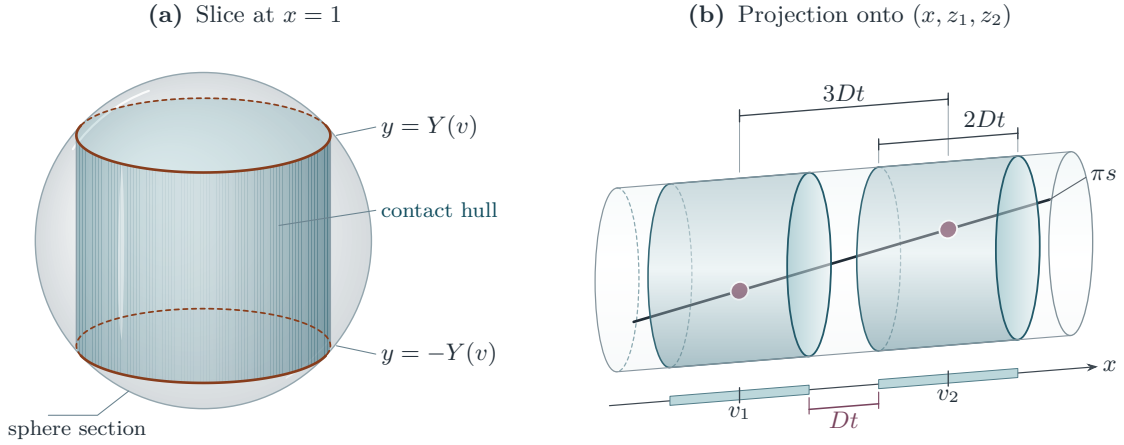


Figure 14: Two views of the elliptic product $\mathcal{E} \times S^1 \subset \mathbb{R}^4$. In (a), the cross-section $x = v = 1$ in coordinates (y, z_1, z_2) has two contact circles at $y = \pm Y(v)$ with $Y(v) = \sqrt{1 - v^2/4}$, whose convex hull fills the shaded solid cylinder inside the section of the supporting sphere. In (b), projecting onto (x, z_1, z_2) preserves axial distances: two marks separated by $3Dt$ have disjoint belts of width $2Dt$ separated by an axial gap of Dt, forcing distinct simplices at distinct successful marks.

We still need to verify that a fixed successful mark q in Proposition 4.1 lies almost surely inside a single simplex. Expose the d sites of any candidate d-tuple successively. Before the k-th site ($1 \leq k \leq d$) is drawn, the affine span of q and the $k-1$ already exposed sites has dimension at most $d-1$. Thus it sits inside an affine hyperplane, which the k-th site misses almost surely by our hyperplane-zero assumption. Hence q and any d sampled sites are affinely independent almost surely. Taking the union over all $\binom{n}{d}$ subsets of d sites excludes q from every facet of every full-dimensional simplex in any sample-vertex refinement; whenever its witness patches are occupied, q lies in the sample hull and therefore inside a unique simplex interior.

Proof of Corollary 4.2. When $\delta = d-1$, the prolate spheroid $\partial E_{R,a}$ of Section 4 satisfies Proposition 4.1. Every interior axial slice is a $(d-2)$-sphere of positive radius whose convex hull fills the cross-section, and all interior levels are noncritical. When $2 \leq \delta \leq d-2$, put $m = d-2 \geq \delta$ and form $V = \mathcal{E} \times M \subset \mathbb{R}^2 \times \mathbb{R}^m = \mathbb{R}^d$, where $\mathcal{E} = \{x^2/4 + y^2 = 1\}$ and $M \subset S^{m-1} \subset \mathbb{R}^m$ is a compact connected real-analytic $(\delta-1)$-manifold of full affine span in $\mathbb{R}^m$. By (4.1), the squared norm on V depends only on x with $h''(x) = 3/2 > 0$. Since $\operatorname{conv}(\mathcal{E} \times M) = (\operatorname{conv}\mathcal{E}) \times (\operatorname{conv} M)$ and M has full affine span on S^{m-1}, the convex hull of each slice $\{\pm\sqrt{1 - v^2/4}\} \times M$ fills $(\operatorname{conv} V) \cap \{x = v\}$ at every $|v| < 2$, while the axial coordinate on $\mathcal{E}$ has critical points only at $x = \pm 2$ (Fig. 14).

It remains to build such a spherical factor $M \subset S^{m-1}$ when $m \geq \delta$. If $m = \delta$, simply take $M = S^{\delta-1}$. If $m > \delta$, send each $u = (u_1, \ldots, u_\delta) \in S^{\delta-1}$ to the unit vector in $\mathbb{R}^m$ proportional to

$$(u_1, \ldots, u_\delta,\ u_1^3,\ u_1^5,\ \ldots,\ u_1^{2(m-\delta)+1}).$$

Normalizing the first δ coordinates recovers u uniquely, so this map is a real-analytic embedding with antipodally symmetric image $M \subset S^{m-1}$. Suppose a linear functional vanished on M, and hence on the displayed polynomial vector. Averaging its values as we flip the signs of $u_2, \ldots, u_\delta$ cancels all terms containing $u_2, \ldots, u_\delta$, leaving an odd polynomial in u_1 that vanishes on

$[-1,1]$ and so vanishes identically. Repeating for each coordinate forces every coefficient to zero. Thus $V = \mathcal{E} \times M$ is connected, real-analytic, without boundary and of full affine span in $\mathbb{R}^d$. By analyticity [26], every affine hyperplane has intrinsic volume zero, and Proposition 4.1 gives the dichotomy between $x(a) = x(b)$ and $x(a) \neq x(b)$ (holding for almost every pair $(a,b) \in (\operatorname{int}\operatorname{conv} V)^2$). □

C The tetrahedra at a fourth-order contact

To prove Theorem 5.1 and Corollary 5.2, we build a space curve in $\mathbb{R}^3$ and two interior endpoint balls U_-, U_+ so that every tetrahedron joining one site near each rear contact R_+, R_- to two sites near the front contact B crosses every segment from U_- to U_+. Subsection C.1 carries out the curve construction via helical arcs and Fourier approximation with exact derivative matching. Subsection C.2 then proves Proposition C.1, tilting the reference sphere to produce an empty circumball through any close rear pair and any consecutive front pair in an $n^{-2/3}$-window. Finally, Subsection C.3 pairs bounded differences over the independent sampled sites with a binomial lower tail on the front windows to establish the simultaneous $n^{4/3-\varepsilon}$ bound, and deduces Corollary 5.2 from Proposition D.2.

C.1 Closing three helical arcs

We start from an exact one-parameter family of spheres and three helical arcs sharing one fourth-order contact and two quadratic contacts. Let H_u denote the screw isometry rotating the first two coordinates by angle u and translating the third by u. Define the base points $B = (1,0,0)$, $R_\pm = (-1, \pm\sqrt{2}, \pm\sqrt{2})$ with matching signs, and centre $C = (-1,0,0)$. Let $P_\tau(p) = |p - H_\tau C|^2 - 4$ be the power of the sphere of radius 2 centred at $H_\tau C$. Applying $H_{-\tau}$ returns the sphere centre to C and sends $H_u X$ to $H_{u-\tau} X$, so $P_\tau(H_u X)$ depends only on $w = u - \tau$. Expanding directly gives

$$\begin{aligned} P_\tau(H_u B) &= w^2 - 2(1 - \cos w) = w^4/12 + O(w^6), \\ P_\tau(H_u R_\pm) &= w^2 + 2(1 - \cos w) \pm 2\sqrt{2}(w - \sin w) = 2w^2 + O(|w|^3). \end{aligned} \tag{C.1}$$

Along the front helix $u \mapsto H_u B$, spheres bitangent at symmetric phases $\pm r$ have centre $(-r/\sin r, 0, 0)$ (studied by Erickson [18, Lemma 2.1] and Bochiş and Santos [5, Lemma 4.2, pp. 116–117]), converging to our limiting sphere centred at C as $r \to 0$.

From these three short helical arcs around $u = 0$, we construct both a C^∞ closed curve keeping the helices exact and a real-analytic closed curve matching all derivatives at the three contacts through order $M \geq 4$. By (C.1), for small $\eta > 0$, all six endpoints of the arcs $|u| \leq 2\eta$ at $\tau = 0$ lie strictly outside the closed ball of centre C and radius 2. Because its exterior in $\mathbb{R}^3$ is path connected, we may join those six endpoints in cyclic order by smooth arcs outside the ball, perturb them slightly so that the connecting arcs are disjoint and meet the helical arcs only at the joins, and smooth the joins. Parametrizing this simple C^∞ loop periodically lets each retained arc keep its screw phase u.

To build the real-analytic loop Γ, fix an integer $M \geq 4$ with $\varepsilon(M-1) > 2$. Truncated Fourier series of the periodic parametrization converge uniformly together with their derivatives up to order $M+1$, though truncation slightly alters derivatives at the three contacts. To restore exact derivatives of orders $0, \dots, M$ at all three contact phases simultaneously, consider the linear map sending a trigonometric polynomial loop to those derivative jets at the three phases. Smooth bump loops supported near each contact realize arbitrary derivative data there, and their Fourier sums approximate them in C^M. Thus the image of this linear map is a dense linear subspace of the finite-dimensional space of derivative data and hence equals the full space. Adding one trigonometric polynomial preimage per basis component supplies

corrections tending to zero in C^{M+1} while matching all derivatives through order M at the three contacts.

Write $\gamma_X(u)$ for the resulting analytic parametrization in local phase u based at $X \in \{B, R_+, R_-\}$. Since $\gamma_X(u)$ and $H_u X$ agree through order M at $u = 0$, Taylor's formula bounds the j-th derivative of $\gamma_X(u) - H_u X$ for $j \in \{0, 1, 2\}$ by $C|u|^{M+1-j}$, with C uniform along the converging Fourier approximations. Pick the three retained arcs small enough that at $\tau = 0$ the leading powers in (C.1) dominate the Taylor errors, and take a high enough Fourier degree. Then Γ stays an embedded trigonometric polynomial circle, $P_0 > 0$ everywhere off the three contact points, and arclength speed stays bounded above and away from zero on each local arc.

To check almost sure general position on Γ, notice that the first three derivatives at B are $(0, 1, 1)$, $(-1, 0, 0)$, and $(0, -1, 0)$, spanning $\mathbb{R}^3$. If a sphere contained Γ, the first three derivatives of its power along Γ would vanish at B, forcing centre C and radius 2; yet for that sphere the fourth derivative at B equals $2 \neq 0$. Every plane and sphere therefore meets Γ in a finite set of arclength zero by analyticity. Exposing sites successively and taking a countable union shows that almost surely no four sites are coplanar and no five are cospherical.

Finally, we place the two interior balls U_-, U_+. In Fig. 9(a), the triangle BR_+R_- lies in the plane $\Pi = \{z = y\}$, which both rear arcs cross transversely because $(z - y)' = 2 \neq 0$ at R_+ and R_-. Any supporting hyperplane to $\operatorname{conv}\Gamma$ through a relative interior point of BR_+R_- would have to contain B, R_+, R_- and equal Π, which is impossible since Γ crosses Π. Thus the barycentre of BR_+R_- lies in $\operatorname{int}\operatorname{conv}\Gamma$. Pick two small open balls $U_-, U_+ \subset \operatorname{int}\operatorname{conv}\Gamma$ on opposite sides of Π along the vertical line through that barycentre, close enough to the triangle that every segment joining U_- to U_+ pierces the relative interior of BR_+R_- away from its edges. That crossing persists under small perturbations of the vertices. Any nondegenerate tetrahedron with one vertex close enough to R_+, one close to R_-, and two close to B therefore has both faces on its rear edge pierced through their relative interiors by every segment joining U_- to U_+, putting an open segment inside the tetrahedron.

C.2 The sphere through four sites

Two phase scales govern the construction: the rear-pair separation $h = 1/n$, and the shrinking matching window $T = n^{-\varepsilon}$ where Γ tracks the helices closely. Take a rear pair near phase t separated by eh $(1 \le e \le 2)$, together with any two front sites in a window of width of order $h^{2/3}$. Proposition C.1 builds a sphere through those four sites whose open ball misses Γ except between the two front sites and within $O(h)$ of the two rear sites.

Proposition C.1. *Fix $0 < \varepsilon < 1/3$ and the curve $\Gamma \subset \mathbb{R}^3$ built in Subsection C.1 with $\varepsilon(M - 1) > 2$. There exist $C > 0$ and an integer n_0 such that the following holds for each $n \ge n_0$. Put $h = 1/n$ and $T = n^{-\varepsilon}$. For any $0 \le t \le T/2$, $1 \le e \le 2$, and $f < g$ in $t + [1, 2]h^{2/3}$, there is a sphere through*

$$\gamma_{R_+}(t), \qquad \gamma_{R_-}(t + eh), \qquad \gamma_B(f), \qquad \gamma_B(g)$$

whose open ball meets Γ only in the union of three sets: the front arc $\gamma_B(u)$ for $f < u < g$, and the rear neighbourhoods $\gamma_{R_+}(u)$ with $|u - t| \le Ch$ and $\gamma_{R_-}(u)$ with $|u - (t + eh)| \le Ch$. The same C and n_0 work uniformly, however small $g - f$.

Proof. Whenever $|u| \le 2T$ and $0 \le \tau \le T$, derivative matching gives $|\gamma_X^{(j)}(u) - H_u^{(j)} X| = O(T^{M+1-j})$ for $j \in \{0, 1, 2\}$. Because $|u| \le 2T$ and $\varepsilon(M - 1) > 2$ makes $T^{M-1} = o(h^2)$, the errors in $P_\tau(\gamma_X(u))$ and its first derivative are $O(T^{M+1})$ and $O(T^M)$, and in its second derivative $O(T^{M-1}) = o(h^2)$. On the rest of the retained arcs where $2T < |u|$, we have

$|u| \le 2|u-\tau|$, so the Taylor errors are absorbed by the leading powers in (C.1). Combining both regimes gives the uniform bounds for $0 \le \tau \le T$:

$$P_\tau(\gamma_{R_\pm}(u)) \ge c|u-\tau|^2 - o(h^2), \qquad (P_\tau \circ \gamma_B)''(u) \ge c|u-\tau|^2 - o(h^2). \tag{C.2}$$

Outside the three retained arcs, compactness and $P_0 > 0$ keep P_τ bounded below by a fixed positive constant for all $0 \le \tau \le T$ and large n.

We build the sphere through the four sites by tuning τ and adding an affine correction to P_τ. Write $\pi(x,y,z) = (x,y)$ for horizontal projection. For $\kappa \in [0,4]$, set $\tau = t + \kappa h$ and let $A_\kappa(x,y)$ be the unique affine function of (x,y) interpolating the values of P_τ at $\gamma_{R_+}(t)$, $\gamma_{R_-}(t+eh)$, and $\gamma_B(f)$. As $n \to \infty$, their horizontal triangle tends to the nondegenerate triangle with vertices $(-1,\sqrt{2})$, $(-1,-\sqrt{2})$, and $(1,0)$. At the rear sites, $P_\tau(\gamma_{R_+}(t)) = 2\kappa^2 h^2 + o(h^2)$ and $P_\tau(\gamma_{R_-}(t+eh)) = 2(e-\kappa)^2 h^2 + o(h^2)$, while at the front site $f - \tau = O(h^{2/3})$ and the jet-matching error above give $P_\tau(\gamma_B(f)) = O(h^{8/3}) + O(T^{M+1}) = o(h^2)$. Thus all coefficients of A_κ are $O(h^2)$. Because both rear projections approach $(-1, \pm\sqrt{2})$ with equal x-coordinates, the y-slope satisfies

$$h^{-2}\partial_y A_\kappa = \frac{2\kappa^2 - 2(e-\kappa)^2}{2\sqrt{2}} + o(1) = \frac{2e\kappa - e^2}{\sqrt{2}} + o(1).$$

Subtracting $A_\kappa \circ \pi$ from P_τ adds only an affine function to $|p|^2$, producing the power of a new sphere through $\gamma_{R_+}(t)$, $\gamma_{R_-}(t+eh)$, and $\gamma_B(f)$. Let $F_\kappa(u) = (P_\tau - A_\kappa \circ \pi)(\gamma_B(u))$ be its restriction to the front arc. Because $\gamma_B'(t)$ has horizontal projection near $(0,1)$, differentiating at $u = t + \xi h^{2/3}$ with $1 \le \xi \le 2$ compares the front slope $w^3/3$ against $\partial_y A_\kappa$:

$$h^{-2}F_\kappa'(u) = \frac{\xi^3}{3} - \frac{2e\kappa - e^2}{\sqrt{2}} + o(1),$$

uniformly over $1 \le \xi \le 2$ and $0 \le \kappa \le 4$. At $\kappa = 0$, this derivative is strictly positive on $t + [1,2]h^{2/3}$ because $\xi^3/3 + e^2/\sqrt{2} \ge 1/3 + 1/\sqrt{2} > 0$. At $\kappa = 4$, it is strictly negative throughout because $\xi^3/3 \le 8/3 < 7/\sqrt{2} \le (8e - e^2)/\sqrt{2}$ for all $1 \le e \le 2$. Since $F_\kappa(f) = 0$ for every κ and $g > f$, integrating gives $F_0(g) > 0 > F_4(g)$, so by continuity some $\kappa \in (0,4)$ achieves $F_\kappa(g) = 0$ as well. Notably, this sign argument works for arbitrarily close front sites $f < g$.

Where can the ball meet the curve? Write $F = F_\kappa$ for the sphere just found. Since $A_\kappa \circ \pi$ has magnitude $O(h^2)$ on compact Γ and $0 \le \tau - t \le 4h$, (C.2) forces positive power everywhere off the three retained arcs and on the rear arcs outside distance Ch of $\gamma_{R_+}(t)$ and $\gamma_{R_-}(t+eh)$. On the front arc, subtracting the affine function $A_\kappa \circ \pi$ adds only a bounded error $O(h^2|\gamma_B''(u)|) = O(h^2)$ to the second derivative, so (C.2) gives $F''(u) \ge c(u-\tau)^2 - Ch^2$. Throughout $[f,g]$, we have $u - \tau \ge h^{2/3} - 4h \ge h^{2/3}/2$ for large n. Thus F is strictly convex on $[f,g]$ with $F(f) = F(g) = 0$, making $F(u) < 0$ on (f,g) and $F'(f) < 0 < F'(g)$.

To rule out any second negative interval on the front arc, expand $F(u)$ around the closer endpoint $v \in \{f,g\}$ by Taylor's integral remainder formula. Setting $\lambda = (u-v)/(v-\tau)$ so that $v - \tau + r(u-v) = (v-\tau)(1+\lambda r)$ for $r \in [0,1]$, we get for every phase u on the front arc

$$F(u) - F'(v)(u-v) = \int_0^1 (1-r)(u-v)^2 F''(v + r(u-v))\, dr \ge (u-v)^2\big(c(v-\tau)^2/6 - Ch^2/2\big),$$

using the elementary bound $\int_0^1 (1-r)(1+\lambda r)^2\, dr = ((\lambda+2)^2+2)/12 \ge 1/6$ for all real λ. Because $|v - \tau| \ge h^{2/3}/2$, the bracket is positive for large n. Applying this with $v = f$ and $F'(f) < 0$ gives $F(u) > 0$ for $u < f$, and applying it with $v = g$ and $F'(g) > 0$ gives $F(u) > 0$ for $u > g$. All front points inside the open ball therefore lie in (f,g). □

C.3 How many such tetrahedra occur?

To convert Proposition C.1 into Theorem 5.1, we lay down a grid of disjoint rear test guards along $[0, T/3]$ and count front sites inside each matching window I_j. Because the n sites are drawn independently and moving one site alters at most two disjoint guard outcomes, bounded differences concentrates the usable rear edges, while a binomial lower tail ensures order $n^{1/3}$ front sites in every window simultaneously.

Choose evenly spaced grid phases $t_j \in [0, T/3]$ at separation Lh for a large constant L fixed below. Around each t_j, define a *rear test* at index j that succeeds if two isolation events occur simultaneously. First, the arc through R_+ receives one site in $t_j + [0, 1/4]h$ and none elsewhere in the guard $t_j + [-D, D]h$; second, the arc through R_- receives one site in $t_j + [3/2, 7/4]h$ and none elsewhere in $t_j + [-D, D]h$. Here we fix D large enough to cover the Ch-neighbourhoods of Proposition C.1, and put $L > 2D$ so distinct guards are pairwise disjoint. Whenever test j succeeds, its rear phases are $t \in t_j + [0, 1/4]h$ and $t + eh$ with $1 \le e \le 2$. Pair each t_j with the deterministic front window $I_j = t_j + [5/4, 7/4]h^{2/3}$ (Fig. 15(a)). Because each target interval has mass at least ch and the union of both guards has mass at most Ch, test j succeeds with probability at least $n(n-1)(ch)^2(1-Ch)^{n-2}$, bounded below by a positive constant independent of n and j.

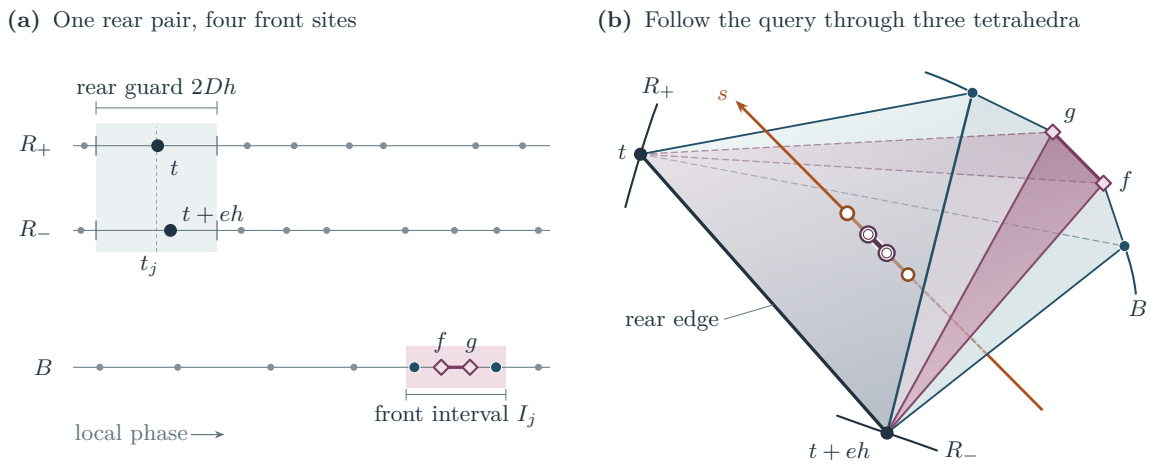


Figure 15: Counting crossed tetrahedra along a shared rear edge. In (a), guard intervals of width $2Dh$ isolate one site on each rear arc, while four front sites inside I_j provide three consecutive front pairs (including f, g). In (b), those three pairs form three Delaunay tetrahedra sharing the same rear edge; along any segment s joining U_- to U_+, four open circles mark triangular face crossings and double rings mark entry into and exit from the middle tetrahedron using f, g (front spacing enlarged).

There are order $T/h = n^{1-\varepsilon}$ disjoint rear tests. Moving one sample site alters at most two guard tests, so McDiarmid's bounded differences inequality [25] ensures at least $cn^{1-\varepsilon}$ successful tests except with probability at most $\exp(-c'n^{1-2\varepsilon})$. Meanwhile, for large n, any rear phase $t \in t_j + [0, 1/4]h$ satisfies $I_j \subset t + [1, 2]h^{2/3}$. Each individual window I_j has sampling mass comparable to $h^{2/3} = n^{-2/3}$, holding a binomial count of order $n^{1/3}$ front sites in expectation. Applying a Chernoff lower-tail bound to each I_j and taking a union bound over all $O(n)$ indices j (which requires no independence across overlapping windows) yields at least $cn^{1/3}$ front sites inside *every* I_j, except with probability $O(n)\exp(-c'n^{1/3})$.

Now pair successful rear edges with consecutive front sites as in Fig. 15(b). Ordering the $m_j \ge 1$ front sites inside I_j by phase gives $m_j - 1$ consecutive pairs, and Proposition C.1 pairs each with the isolated rear edge of test j to form a tetrahedron with empty open circumball. By Subsection C.1, sample general position holds almost surely, making each tetrahedron a full-dimensional Delaunay cell pierced by every segment from U_- to U_+. Crucially, even if

tests share front sites, their rear vertices sit in disjoint guards, keeping all counted tetrahedra distinct. Let G_n be the event where both concentration bounds and general position hold. On G_n, simultaneously for *every* segment s joining U_- to U_+,

$$N_n(s) \geq \sum_{j \text{ successful}} (m_j - 1) \geq c\, n^{1-\varepsilon} n^{1/3}. \tag{C.3}$$

Because $0 < \varepsilon < 1/3$ implies $1 - 2\varepsilon > 1/3$, combining both failure bounds gives $\mathbb{P}(G_n^c) \leq \exp(-c_1 n^{1/3})$ for some $c_1 > 0$ and large n. By Borel–Cantelli, (C.3) holds eventually almost surely simultaneously for all joining segments (even chosen after sampling), and taking expectations gives $\mathbb{E}N_n(s) \geq cn^{4/3-\varepsilon}\mathbb{P}(G_n)$ for Theorem 5.1.

Proof of Corollary 5.2. Take the compact C^∞ embedded loop Γ_∞ from Subsection C.1 prior to Fourier approximation, which contains the exact helical arcs $b(u) = H_u B$ and $r_\pm(u) = H_u R_\pm$ over an open phase interval around 0. Each helix is injective (its third coordinate grows at unit speed) with positive constant speed. Fix $\eta > 0$ small enough that the three arcs are exact and pairwise disjoint for $|u| \leq 2\eta$, and that by (C.1) the smooth quotients $P_t(b(u))/(u-t)^4$ and $P_t(r_\pm(u))/(u-t)^2$ (which depend only on $u - t$ and extend through zero with values $1/12, 2, 2$) stay strictly positive on $|u - t| \leq 3\eta$.

At $t = 0$, the compact complement of the open arcs $|u| < 2\eta$ lies outside the reference sphere ($P_0 > 0$). By continuity, $P_t > 0$ persists on that complement for all $t \in [0, T_0]$ once $0 < T_0 < \eta$ is chosen small. Restricting the parametrizations to $J = [-\eta, T_0 + \eta]$, any helical point with $|u| \leq 2\eta$ outside J satisfies $\eta \leq |u - t| \leq 3\eta$ for $t \in [0, T_0]$, keeping $P_t > 0$ there as well. Thus Γ_∞ satisfies (D.1) and all positivity hypotheses of Appendix D.1 with constant arclength density, and Proposition D.2 gives U_-, U_+ and the $n^{4/3}$ bounds. □

D An analytic curve with two components

In Appendix C, the single real-analytic curve approximates a fourth-order contact only on a shrinking interval, which loses a factor n^ε in the count of rear edges. Here we first prove a general lower bound for any smooth curve (possibly with multiple components) that supports a one-parameter family of empty spheres with one fourth-order contact and two quadratic contacts over a fixed interval. This single proposition yields both the C^∞ bound of Corollary 5.2 and the two-component real-analytic construction of Corollary 5.3. We then realize those exact contact orders explicitly on two stereographically projected circles in $\mathbb{R}^3$.

D.1 A quartic contact along a fixed interval

In Appendix C, the single analytic curve matched the helical contact orders at one point and lost a factor n^ε in the count of rear edges. Whenever an exact one-parameter family of empty spheres keeps one fourth-order contact and two quadratic contacts over a *fixed* interval $[0, T_0]$, we adapt the empty-circumball construction of Subsection C.2 and the disjoint-guard count of Subsection C.3. This supplies order n rear edges and reaches the full exponent $4/3$, whether the curve has one component (Corollary 5.2) or two (Corollary 5.3).

Let $\Gamma \subset \mathbb{R}^3$ be a compact smooth embedded curve without boundary, possibly with multiple connected components. Choose three pairwise disjoint arcs on Γ with smooth injective parametrizations $b, r_+, r_- \colon J \to \Gamma$ of nonzero velocity, defined on $J = [-\eta, T_0 + \eta]$ for fixed constants $T_0, \eta > 0$. Here $u \in J$ is the phase along each arc, with b tracing the front arc and r_+, r_- tracing the two rear arcs. Suppose there is a smooth one-parameter family of spheres indexed by $t \in [0, T_0]$, with centre $c(t)$, radius $\rho(t)$ and power $P_t(x) = |x - c(t)|^2 - \rho(t)^2$. Assume that along the three arcs,

$$P_t(b(u)) = a_B(u,t)(u-t)^4, \qquad P_t(r_\pm(u)) = a_\pm(u,t)(u-t)^2, \tag{D.1}$$

where a_B, a_+, a_- are continuous and strictly positive on $J \times [0, T_0]$, while $P_t > 0$ everywhere else on $\Gamma \setminus \{b(t), r_+(t), r_-(t)\}$.

Before constructing the empty circumballs and counting tetrahedra in Proposition D.2, we check that these contact orders force the three contacts onto the osculating plane of b. Lemma D.1 proves that the three contacts lie on the osculating circle of b at $b(t)$, that the relative interior of their triangle lies in $\operatorname{int}\operatorname{conv}\Gamma$, and that the rear chord coefficients have opposite signs, which allows us to tilt the sphere across the front window.

Lemma D.1. *Let the three arcs of $\Gamma \subset \mathbb{R}^3$ and the family of spheres satisfy* (D.1) *with the positivity conditions above. For every $t \in [0, T_0]$, the contacts $b(t), r_+(t), r_-(t)$ lie on the osculating circle of b at $b(t)$, and the relative interior of their triangle lies in* $\operatorname{int}\operatorname{conv}\Gamma$. *Moreover, after interchanging the two rear labels once if necessary, the coefficients in the unique decomposition*

$$b'(t) = \alpha_+(t)(r_+(t) - b(t)) + \alpha_-(t)(r_-(t) - b(t))$$

are continuous and satisfy $\alpha_+(t) > 0 > \alpha_-(t)$ throughout $[0, T_0]$.

Proof. At any fixed $t \in [0, T_0]$, consider the spatial affine function $\Lambda_t(x) = \partial_t P_t(x)$, which holds x fixed while differentiating P_t in t, and has linear part $-2c'(t)$. Whenever $x(t)$ moves along a contact of P_t, differentiating $P_t(x(t)) = 0$ gives $\Lambda_t(x(t)) + \nabla_x P_t(x(t)) \cdot x'(t) = 0$. Since $x'(t)$ is tangent to the sphere, the second term vanishes, leaving $\Lambda_t(x(t)) = 0$ at all three contacts. Put $F(u, t) = P_t(b(u))$. Along the diagonal $u = t$, the u-derivatives of orders 0 through 3 vanish identically while $F_{uuuu}(t, t) = 24a_B(t, t) > 0$. Differentiating along $u = t$ therefore yields

$$\frac{d^j}{du^j}\Lambda_t(b(u))\big|_{u=t} = 0 \quad (0 \le j \le 2), \qquad \frac{d^3}{du^3}\Lambda_t(b(u))\big|_{u=t} = -24a_B(t, t) \neq 0.$$

Hence $c'(t) \neq 0$, making $\Pi_t = \{\Lambda_t = 0\}$ a well-defined plane containing $b(t), r_+(t), r_-(t)$. Moreover, the vanishing first two derivatives place both $b'(t)$ and $b''(t)$ parallel to Π_t, while the nonzero third derivative implies that the front arc b crosses Π_t with third-order contact at $u = t$. Next, $F_u(t, t) = 0$ and $F_{uu}(t, t) = 0$ give $(b(t) - c(t)) \cdot b'(t) = 0$ and $(b(t) - c(t)) \cdot b''(t) = -|b'(t)|^2$. Since $b'(t) \neq 0$, the vectors $b'(t)$ and $b''(t)$ are linearly independent and span Π_t, identifying Π_t as the osculating plane of b at $b(t)$. Projecting $c(t)$ orthogonally onto Π_t preserves both dot products, showing that the projected centre is the centre of curvature and hence that the sphere meets Π_t along the osculating circle of b at $b(t)$.

Because the three contacts $b(t), r_+(t), r_-(t)$ are distinct points on this circle, the rear chords $r_+(t) - b(t)$ and $r_-(t) - b(t)$ form a basis for the directions parallel to Π_t. Inside Π_t, let $N(t)$ be the unit inward normal to the osculating circle at $b(t)$, which is orthogonal to $b'(t)$. Both chords point strictly into the half-plane containing the interior of the circle and therefore have strictly positive inner product with $N(t)$. Taking the inner product of $b'(t) = \alpha_+(t)(r_+(t) - b(t)) + \alpha_-(t)(r_-(t) - b(t))$ with $N(t)$ shows that the two positive terms would have to add to zero unless $\alpha_+(t)$ and $\alpha_-(t)$ are nonzero (since $b'(t) \neq 0$) and carry opposite signs. Continuity over $[0, T_0]$ preserves this sign separation after at most one relabelling of r_+ and r_-. Finally, any supporting plane to $\operatorname{conv}\Gamma$ through a relative interior point of the contact triangle would contain all three vertices and thus coincide with Π_t, which is ruled out because b crosses Π_t with third-order contact at $b(t)$. □

Proposition D.2. *Let $\Gamma \subset \mathbb{R}^3$ be a compact smooth embedded curve without boundary whose three arcs and supporting spheres satisfy* (D.1) *and the positivity conditions above. Then there exist open balls U_-, U_+ with disjoint closures in* $\operatorname{int}\operatorname{conv}\Gamma$ *having the following property. Draw a sequence of independent sites on Γ from a law assigning zero probability to sets of arclength zero. On the three marked arcs, assume that every measurable subset occupying a*

fraction r of the total length of Γ receives probability between mr and Mr for fixed constants $0 < m \le M < \infty$. Then there exist $c, c_1 > 0$ and an integer n_0 such that for every $n \ge n_0$, with probability at least $1 - \exp(-c_1 n^{1/3})$, at least $cn^{4/3}$ tetrahedra belong to every Delaunay triangulation of the first n sites and have interiors meeting every segment joining U_- to U_+. Almost surely this holds for all sufficiently large n, and every fixed joining segment s satisfies $\liminf_{n\to\infty} n^{-4/3}\mathbb{E}N_n(s) \ge c$.

Proof. Pick the centres of two small open balls $U_-, U_+ \subset \operatorname{int}\operatorname{conv}\Gamma$ on opposite sides of the triangle $b(0)r_+(0)r_-(0)$ along the normal through its barycentre. By Lemma D.1 and compactness, shrinking both balls ensures that every segment joining U_- to U_+ pierces $b(0)r_+(0)r_-(0)$ with a uniform margin from its edges. Picking $0 < T \le T_0$ sufficiently small preserves this crossing for every triangle formed on the three arcs at phases in $[0, T]$. Any nondegenerate tetrahedron built from two front sites and one site on each rear arc at phases in $[0, T]$ therefore has both faces on its rear edge pierced through their relative interiors by every segment joining U_- to U_+, placing an open subsegment inside the tetrahedron.

Finding the circumsphere. Put $h = 1/n$, choose rear sites $r_+(t)$ and $r_-(t + eh)$ with $0 \le t \le T/2$ and $1 \le e \le 2$, and choose front sites $b(f)$ and $b(g)$ with $f < g$ in $t + [\xi_1, \xi_2]h^{2/3}$, where $0 < \xi_1 < \xi_2$ will be fixed independently of n. For any $k \in [0, e]$, put $\tau = t + kh$. Define A_k on $\mathbb{R}^3$ by interpolating the values of P_τ at $r_+(t)$, $r_-(t + eh)$, and $b(f)$ on their affine plane and extending constantly along its normal. At these three sites, P_τ equals $a_+(t, \tau)k^2h^2$, $a_-(t + eh, \tau)(e - k)^2h^2$, and $O(h^{8/3})$, and their triangle converges uniformly to the nondegenerate contact triangle at phase t. Hence both A_k and ∇A_k are $O(h^2)$ on compact Γ, making $Q_k = P_\tau - A_k$ the power of a sphere through $r_+(t)$, $r_-(t + eh)$, and $b(f)$.

To vary Q_k across the front window, extend the three barycentric coordinates of the interpolation triangle constantly along its normal. Their gradients converge uniformly to those of the contact triangle at phase t, so the directional derivatives along $b'(t)$ of the two rear barycentric coordinates approach the coefficients $\alpha_+(t)$ and $\alpha_-(t)$ from Lemma D.1. Put $u = t + \xi h^{2/3}$ and $D_\pm(t) = a_\pm(t, t)\alpha_\pm(t)$, where Lemma D.1 ensures $D_+(t) > 0 > D_-(t)$ on $[0, T]$. Differentiating $F(u, \tau) = P_\tau(b(u))$ yields $F_u(u, \tau) = 4a_B(\tau, \tau)(u - \tau)^3 + O((u - \tau)^4)$. Subtracting $\nabla A_k \cdot b'(u)$ leaves

$$h^{-2}(Q_k \circ b)'(u) = 4a_B(t, t)\xi^3 - D_+(t)k^2 - D_-(t)(e - k)^2 + o(1)$$

uniformly over $t \in [0, T/2]$, $e \in [1, 2]$, $k \in [0, e]$, and $\xi \in [\xi_1, \xi_2]$. Fix $\xi_2 > 0$ small enough that $4 \max_t a_B(t, t)\xi_2^3 < \min_t D_+(t)/2$ and set $\xi_1 = \xi_2/2$. Since $D_-(t) < 0 < D_+(t)$ and $e \ge 1$, at $k = 0$ the term $-D_-(t)e^2 > 0$ makes the derivative positive throughout the front window. At $k = e$, our choice of ξ_2 keeps $4a_B(t, t)\xi^3 - D_+(t)e^2 < -\frac{1}{2}D_+(t) < 0$ across the window. Because every Q_k vanishes at $b(f)$ and $g > f$, integrating gives $Q_0(b(g)) > 0 > Q_e(b(g))$ regardless of how small $g - f$ is, yielding $k \in (0, e)$ with $Q = Q_k$ vanishing at all four sites. Moreover, the four sites are affinely independent: the affine difference $Q_e - Q_0$ vanishes on the plane of the first three sites, whereas $Q_e(b(g)) < 0 < Q_0(b(g))$ keeps $b(g)$ off that plane.

Locating the ball. On both rear arcs, uniform positivity of $a_\pm$ and the $O(h^2)$ size of A_k give $Q(r_\pm(u)) \ge a_0(u - \tau)^2 - Ch^2$ for a uniform $a_0 > 0$, forcing $Q > 0$ outside $|u - t| \le D_0 h$ for a fixed D_0. On the compact closure of the complement of the three arcs, P_τ stays bounded away from zero (vanishing only inside the open interior of J), so the $O(h^2)$ shift keeps $Q > 0$ there as well. On the front arc, put $H(u) = Q(b(u))$. In a fixed neighbourhood of τ, uniform Taylor expansion gives $H''(u) \ge a(u - \tau)^2 - Ch^2$ with $a > 0$, while compactness keeps $H > 0$ elsewhere on b. On $[f, g]$, we have $u - \tau \ge \xi_1 h^{2/3}/2$ for large n, so H is strictly convex between its roots f and g, giving $H'(f) < 0 < H'(g)$. For any u in the Taylor neighbourhood outside

$[f,g]$, expanding around the closer endpoint $v \in \{f,g\}$ with $\lambda = (u-v)/(v-\tau)$ again yields

$$H(u) \geq H'(v)(u-v) + (u-v)^2\big(a(v-\tau)^2/6 - Ch^2/2\big) > 0$$

for large n, since $H'(v)(u-v) > 0$ and $|v-\tau| \geq \xi_1 h^{2/3}/2$. The closed ball of Q thus meets Γ solely inside $[f,g]$ on the front arc and within $|u-t| \leq D_0 h$ on each rear arc.

Counting the tetrahedra. Lay down the disjoint guard grid of Subsection C.3 on $[0,T/3]$. At phases t_j spaced by Lh, require one site on r_+ in $t_j+[0,1/4]h$, one site on r_- in $t_j+[3/2,7/4]h$, and no other rear sites in either guard $t_j + [-D,D]h$, choosing $D > D_0 + 2$ and $L > 2D$. Because nonzero-velocity arcs convert phase length to arclength up to uniform factors, each target interval has mass at least $a'h$ and the two guards together have mass at most $A'h$. Every test therefore succeeds with probability at least $n(n-1)(a'h)^2(1-A'h)^{n-2} \geq c_0 > 0$. Across the order n pairwise disjoint guards, McDiarmid's bounded differences inequality yields at least $c'n$ successful rear tests outside probability $\exp(-c''n)$. Fixing $\xi_1 < \xi_1' < \xi_2' < \xi_2$, every interval $t_j + [\xi_1', \xi_2']h^{2/3}$ lies inside $t + [\xi_1,\xi_2]h^{2/3}$ for large n and has mass of order $h^{2/3}$. A binomial lower tail and union bound thus give at least $c'n^{1/3}$ front sites in every window except with probability $O(n)\exp(-c''n^{1/3})$. Pairing each successful rear edge with every consecutive pair of front sites in its window produces distinct tetrahedra whose closed circumballs contain no other sampled site. At least $cn^{4/3}$ tetrahedra therefore belong to *every* Delaunay triangulation and meet every segment from U_- to U_+. Since $\exp(-c''n) + O(n)\exp(-c''n^{1/3}) \leq \exp(-c_1 n^{1/3})$ is summable in n, Borel–Cantelli gives the eventual almost-sure bound for all joining segments at once, and taking expectations yields the fixed-segment liminf. □

D.2 Two analytic components on one family of supporting spheres

We now construct the two disjoint real-analytic circles of Corollary 5.3 explicitly. Recall that stereographic projection $\sigma\colon S^3 \setminus \{(0,0,1,0)\} \to \mathbb{R}^3$ from the pole $(0,0,1,0) \in S^3 \subset \mathbb{R}^4$ onto the equatorial hyperplane $x_3 = 0$ maps every 2-sphere section of S^3 missing the pole onto a sphere in $\mathbb{R}^3$ (and maps the open spherical cap containing the pole, minus the pole itself, to the sphere's exterior). Put two closed curves on S^3 modeled on the trigonometric moment curve $u \mapsto (\cos u, \sin u, \cos 2u, \sin 2u)$ [30, Sec. 5, Eq. (5.1)], and sweep a moving hyperplane across $\mathbb{R}^4$. We make the hyperplane height a trigonometric polynomial with one fourth-order zero on the first curve and two quadratic zeros on the second, realizing Proposition D.2 upon projection.

Proof of Corollary 5.3. Identify $\mathbb{R}^4$ with $\mathbb{C}^2$ via $z_1 = x_1 + ix_2$ and $z_2 = x_3 + ix_4$, so S^3 is given by $|z_1|^2 + |z_2|^2 = 1$. Put $a = \sqrt{14/15}$ and $b = 1/\sqrt{15}$, and define two parametrized curves on S^3 for $u \in \mathbb{R}/(2\pi\mathbb{Z})$ by

$$X(u) = (ae^{iu}, be^{2iu}), \qquad Y(u) = (\tfrac{a}{4}(1+3i)e^{iu}, \tfrac{b}{2}(-4+3i)e^{2iu}).$$

One checks directly that $a^2 + b^2 = 1$ and $\frac{10}{16}a^2 + \frac{25}{4}b^2 = 1$, so both curves lie on S^3. On each curve, z_1 traces a circle once at constant nonzero speed, making both parametrizations injective and regular. Because those two circles in $\mathbb{C}$ have distinct radii a and $\frac{\sqrt{10}}{4}a$, the curves X and Y are disjoint, and neither reaches the projection pole $(0,0,1,0)$ where $z_1 = 0$. Stereographic projection $\sigma(x) = (x_1, x_2, x_4)/(1-x_3)$ therefore produces two disjoint real-analytic embedded circles $\beta = \sigma \circ X$ and $\varrho = \sigma \circ Y$ in $\mathbb{R}^3$.

Next, for $\tau \in [0,1]$, consider the moving affine functional on $\mathbb{R}^4$ given by

$$f_\tau(x) = 3 - \tfrac{4}{a}(x_1\cos\tau + x_2\sin\tau) + \tfrac{1}{b}(x_3\cos 2\tau + x_4 \sin 2\tau).$$

Set $w = u - \tau$ and $u_1 = \pi + \arctan(3/4)$, so that $\cos u_1 = -4/5$ and $\sin u_1 = -3/5$. Substituting $X(u)$ and $Y(u)$ into f_τ yields the exact trigonometric factorizations

$$f_\tau(X(u)) = 3 - 4\cos w + \cos 2w = 2(1 - \cos w)^2, \tag{D.2}$$

$$f_\tau(Y(u)) = (1 - \cos w)(5 + 4\cos w + 3\sin w) = 5(1 - \cos w)(1 - \cos(w - u_1)). \tag{D.3}$$

Because $1 - \cos w = \frac{1}{2}w^2 + O(w^4)$, equation (D.2) has one fourth-order zero at $u = \tau$ modulo 2π, while (D.3) has two quadratic zeros at $u = \tau$ and $u = \tau + u_1$ modulo 2π. Both functions stay strictly positive everywhere else.

To translate these nonnegative hyperplane heights on S^3 into powers of empty balls in $\mathbb{R}^3$, put $y = \sigma(x) \in \mathbb{R}^3$ and invert σ:

$$x = \frac{(2y_1, 2y_2, |y|^2 - 1, 2y_3)}{1 + |y|^2}.$$

Clearing denominators shows that $(1 + |y|^2)f_\tau(x)$ is a quadratic polynomial in $y \in \mathbb{R}^3$ with leading term $m_\tau|y|^2$, where $m_\tau = 3 + \sqrt{15}\cos 2\tau$. Since $2\tau \in [0, 2]$ keeps $\cos 2\tau \geq \cos 2 > -0.42$, we have $m_\tau > 1.37 > 0$ throughout $\tau \in [0, 1]$. Dividing $(1 + |y|^2)f_\tau(x)$ by m_τ and completing the square yields the power $P_\tau(y) = |y - c(\tau)|^2 - \rho(\tau)^2$ of a sphere $S_\tau \subset \mathbb{R}^3$. Its radius is strictly positive because its zero set already contains three distinct contacts. Because $f_\tau \geq 0$ on both curves, β and ϱ miss the open ball of S_τ, and the sphere power $P_\tau(y)$ equals (D.2) or (D.3) times the smooth positive factor $(1 + |y|^2)/m_\tau$. Figure 16(a) draws the two linked circles and S_0, while Fig. 16(b) displays the quartic and quadratic contact orders.

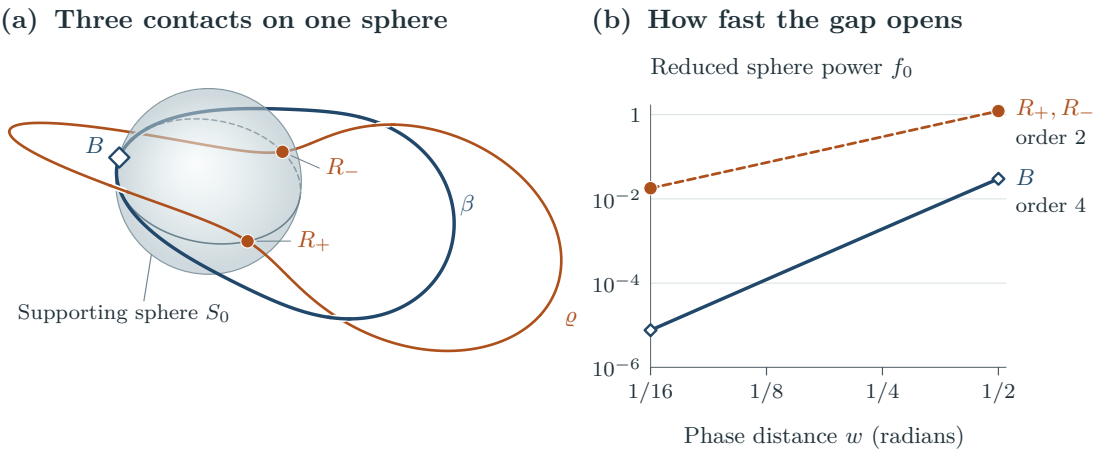


Figure 16: One sphere and three contacts on two disjoint real-analytic circles. In (a), the sphere S_0 touches β at the diamond B and touches ϱ at the two dots R_+, R_-; the gray circle through all three contacts osculates β at B. In (b), on logarithmic axes, the curves plot the normalized power f_0 from (D.2) and (D.3) (which divides the sphere power by $(1 + |y|^2)/m_0$). At phase distance $w > 0$ from each contact, f_0 grows asymptotically as $w^4/2$ at B and $9w^2/2$ at either rear contact, with both rear curves coinciding because the two factors in (D.3) swap roles between phases w and $u_1 - w$.

Now restrict the three arcs to the fixed interval $J = [-1/10, 11/10]$ via $b(u) = \beta(u)$, $r_+(u) = \varrho(u)$, and $r_-(u) = \varrho(u + u_1)$. Because $u_1 = \pi + \arctan(3/4) \approx 3.785$, the intervals J (of length 1.2) and $J + u_1$ are disjoint modulo 2π. Along $r_+(u) = \varrho(u)$, the phase offset in (D.3) is $w = u - \tau$, so the first factor $1 - \cos(u - \tau)$ vanishes at $u = \tau$ while the second factor is $1 - \cos(u - \tau - u_1)$. Along $r_-(u) = \varrho(u + u_1)$, the offset is $w = u - \tau + u_1$, so the two factors swap roles: the second factor $1 - \cos(u - \tau)$ vanishes at $u = \tau$ while the first is $1 - \cos(u - \tau + u_1)$. For every $u \in J$ and $\tau \in [0, 1]$, we have $|u - \tau| \leq 11/10 < 2\pi - u_1 < u_1$, which keeps $u - \tau - u_1 \in (-2\pi, 0)$ and $u - \tau + u_1 \in (0, 2\pi)$; thus on each rear arc, the

factor associated with the other rear contact stays strictly positive on $J \times [0,1]$ and equals $1 - \cos u_1 = 9/5$ at $u = \tau$. Because $(1 - \cos z)/z^2$ extends smoothly through $z = 0$ with value $1/2$, the reduced power f_τ has leading quadratic term $5(9/5)(u-\tau)^2/2 = \frac{9}{2}(u-\tau)^2$ on both rear arcs, and dividing P_τ by $(u-\tau)^4$ on b and by $(u-\tau)^2$ on $r_\pm$ gives continuous strictly positive coefficients a_B, a_+, a_- on $J \times [0,1]$. Off the three arcs (and at their endpoints), compactness keeps P_τ uniformly positive. Because both circles have nonzero speed bounded above and away from zero, normalized arclength on $\beta \cup \varrho$ satisfies all density bounds of Proposition D.2, supplying U_-, U_+ and the $n^{4/3}$ lower bounds.

Finally, verify almost sure general position on $\beta \cup \varrho$. Pullbacks of any plane or sphere in $\mathbb{R}^3$ under σ give affine hyperplane equations in $\mathbb{R}^4$. Since $1, \cos u, \sin u, \cos 2u, \sin 2u$ are linearly independent and both complex amplitudes of X and Y are nonzero, those hyperplanes restrict along $X(u)$ and $Y(u)$ to nonzero trigonometric polynomials of degree at most two. Every plane, sphere, or line in $\mathbb{R}^3$ therefore meets $\beta \cup \varrho$ in finitely many points, so conditioning on previously sampled sites excludes collinear triples, coplanar quadruples, and cospherical quintuples almost surely. □

Data and code availability

The research archive contains the sampling seeds, query specifications and per-sample measurements for the comparisons of dimension and segment length. For the two space curves, it identifies the rear edge and the four sites of each marked tetrahedron visited by the segment. The code reproduces the estimates and figures for all three experiments from these data. The author will release these materials in a public repository once the preprint is public on both arXiv and Hexagon. The SHA-256 checksum of the release archive is `03b1b92bb4d173858d55cdd84836870af46e667e496744b1201128bc0302ab77`

References

[1] Nina Amenta, Dominique Attali, and Olivier Devillers. A tight bound for the Delaunay triangulation of points on a polyhedron. *Discrete & Computational Geometry*, 48(1):19–38, 2012. doi:10.1007/s00454-012-9415-7.

[2] Nina Amenta and Günter M. Ziegler. Deformed products and maximal shadows of polytopes. In *Advances in Discrete and Computational Geometry*, volume 223 of *Contemporary Mathematics*, pages 57–90. American Mathematical Society, 1999. doi:10.1090/conm/223/03132.

[3] Charles Arnal, David Cohen-Steiner, and Vincent Divol. Critical points of the distance function to a generic submanifold. *Algebraic & Geometric Topology*, 26(8):2915–2986, 2026. doi:10.2140/agt.2026.26.2915.

[4] Dominique Attali, Jean-Daniel Boissonnat, and André Lieutier. Complexity of the Delaunay triangulation of points on surfaces: The smooth case. In *Proceedings of the Nineteenth Annual Symposium on Computational Geometry*, pages 201–210. ACM, 2003. doi:10.1145/777792.777823.

[5] Daciana Bochiş and Francisco Santos. On the number of facets of three-dimensional Dirichlet stereohedra II: Non-cubic groups. *Beiträge zur Algebra und Geometrie*, 47(1):89–120, 2006. URL: https://ftp.gwdg.de/pub/misc/EMIS/journals/BAG/vol.47/no.1/b47h1san.pdf.

[6] Prosenjit Bose and Luc Devroye. Intersections with random geometric objects. *Computational Geometry*, 10:139–154, 1998. doi:10.1016/S0925-7721(98)00004-2.

[7] Prosenjit Bose and Luc Devroye. On the stabbing number of a random Delaunay triangulation. *Computational Geometry*, 36(2):89–105, 2007. doi:10.1016/j.comgeo.2006.05.005.

[8] Madeline Brandt and Madeleine Weinstein. Voronoi cells in metric algebraic geometry of plane curves. *Mathematica Scandinavica*, 130(1), 2024. arXiv:1906.11337v3. doi:10.7146/math.scand.a-139875.

[9] Nicolas Chenavier and Olivier Devillers. Stretch factor of long paths in a planar Poisson-Delaunay triangulation, 2016. arXiv:1607.05770v2. Related journal article: *Stretch factor in a planar Poisson–Delaunay triangulation with a large intensity*, Advances in Applied Probability 50 (2018), 35–56, https://doi.org/10.1017/apr.2018.3. URL: https://arxiv.org/abs/1607.05770v2.

[10] Pedro Machado Manhães de Castro. *Practical Ways to Accelerate Delaunay Triangulations*. PhD thesis, Université Nice Sophia Antipolis, 2010. URL: https://theses.hal.science/tel-00531765.

[11] Pedro Machado Manhães de Castro and Olivier Devillers. Practical distribution-sensitive point location in triangulations. *Computer Aided Geometric Design*, 30(5):431–450, 2013. doi:10.1016/j.cagd.2013.02.004.

[12] Olivier Devillers and Charles Duménil. A Poisson sample of a smooth surface is a good sample. 35th European Workshop on Computational Geometry, 2019. URL: https://www.eurocg2019.uu.nl/papers/40.pdf.

[13] Olivier Devillers, Jeff Erickson, and Xavier Goaoc. Empty-ellipse graphs. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1249–1257, 2008. URL: https://jeffe.cs.illinois.edu/pubs/eellipse.html.

[14] Olivier Devillers, Marc Glisse, Xavier Goaoc, and Rémy Thomasse. Smoothed complexity of convex hulls by witnesses and collectors. *Journal of Computational Geometry*, 7(2):101–144, 2016. doi:10.20382/jocg.v7i2a6.

[15] Luc Devroye, Christophe Lemaire, and Jean-Michel Moreau. Expected time analysis for Delaunay point location. *Computational Geometry*, 29(2):61–89, 2004. URL: https://luc.devroye.org/Devroye+Lemaire+Moreau-DelaunayPointLocation-2004.pdf, doi:10.1016/j.comgeo.2004.02.002.

[16] Luc Devroye, Ernst P. Mücke, and Binhai Zhu. A note on point location in Delaunay triangulations of random points. *Algorithmica*, 22(4):477–482, 1998. doi:10.1007/PL00009234.

[17] Charles Duménil. *Expected size of the 3-dimensional Delaunay triangulation of random points on a surface*. PhD thesis, Université de Lorraine, 2022. doi:10.70675/8bd9f0ffz42b9z421fz995az91a273dfd224.

[18] Jeff Erickson. Nice point sets can have nasty Delaunay triangulations. *Discrete & Computational Geometry*, 30(1):109–132, 2003. doi:10.1007/s00454-003-2927-4.

[19] Martin Golubitsky and Victor Guillemin. *Stable Mappings and Their Singularities*, volume 14 of *Graduate Texts in Mathematics*. Springer, 1973. doi:10.1007/978-1-4615-7904-5.

[20] Ross Hemsley. *Probabilistic Methods for the Analysis of Algorithms on Random Tessellations*. PhD thesis, Université Nice Sophia Antipolis, 2014. URL: https://inria.hal.science/tel-01099165v2.

[21] Qi Huangfu and J. A. J. Hall. Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, 10(1):119–142, 2018. doi:10.1007/s12532-017-0130-5.

[22] Daniel Hug and Rolf Schneider. Large cells in Poisson–Delaunay tessellations. *Discrete & Computational Geometry*, 31:503–514, 2004. doi:10.1007/s00454-003-0818-3.

[23] Clément Jamin, Sylvain Pion, and Monique Teillaud. 3D triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 6.1.1 edition, 2026. URL: https://doc.cgal.org/6.1.1/Triangulation_3/index.html.

[24] Alexander Magazinov. Asymptotics for some combinatorial characteristics of the convex hull of a Poisson point process in the Clifford torus. *Discrete & Computational Geometry*, 49(2):200–220, 2013. doi:10.1007/s00454-012-9476-7.

[25] Colin McDiarmid. On the method of bounded differences. In J. Siemons, editor, *Surveys in Combinatorics, 1989*, pages 148–188. Cambridge University Press, 1989. doi:10.1017/CBO9781107359949.008.

[26] Boris S. Mityagin. The zero set of a real analytic function. *Mathematical Notes*, 107:529–530, 2020. doi:10.1134/S0001434620030189.

[27] Ernst P. Mücke, Isaac Saias, and Binhai Zhu. Fast randomized point location without preprocessing in two- and three-dimensional Delaunay triangulations. *Computational Geometry*, 12(1–2):63–83, 1999. doi:10.1016/s0925-7721(98)00035-2.

[28] Leandro P. R. Pimentel and Raphaël Rossignol. Greedy polyominoes and first-passage times on random Voronoi tilings. *Electronic Journal of Probability*, 17(12):1–31, 2012. doi:10.1214/EJP.v17-1788.

[29] Kristian Ranestad and Bernd Sturmfels. The convex hull of a variety. In Petter Brändén, Mikael Passare, and Mihai Putinar, editors, *Notions of Positivity and the Geometry of Polynomials*, Trends in Mathematics, pages 331–344. Springer Basel, 2011. doi:10.1007/978-3-0348-0142-3_18.

[30] Raman Sanyal, Frank Sottile, and Bernd Sturmfels. Orbitopes. *Mathematika*, 57(2):275–314, 2011. doi:10.1112/S002557931100132X.

[31] Rolf Schneider and Wolfgang Weil. *Stochastic and Integral Geometry*. Probability and Its Applications. Springer, 2008. doi:10.1007/978-3-540-78859-1.

[32] Jonathan Richard Shewchuk. Stabbing Delaunay tetrahedralizations. *Discrete & Computational Geometry*, 32(3):339–343, 2004. doi:10.1007/s00454-004-1095-5.