Geometric Sampling Theory, Triangulations, and Robust Machine Learning

by

Marc Kzhaya Khoury

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Jonathan Richard Shewchuk, Chair
Professor Benjamin Recht
Professor Ming Gu
Dr. Valdlen Koltun

Spring 2020

Geometric Sampling Theory, Triangulations, and Robust Machine Learning

Abstract

Geometric Sampling Theory, Triangulations, and Robust Machine Learning

by

Marc Kzhaya Khoury

Doctor of Philosophy in Computer Science

University of California, Berkeley

Professor Jonathan Richard Shewchuk, Chair

Manifold triangulation is the problem of, given a set of points V sampled densely from some known manifold Σ, output a set of simplices $\mathcal{T}$ that is topologically identical to Σ and geometrically close to Σ. In the first part of this thesis we study a variant of this problem with one additional constraint. In addition to a set of points V densely sampled from a known Σ we are also given a finite set of line segments S, whose endpoints are in V, which must appear as edges in the output triangulation $\mathcal{T}$. To solve this problem we introduce *restricted constrained Delaunay triangulations* (restricted CDTs), which combine ideas from constrained Delaunay triangulations and restricted Delaunay triangulations to enable the enforcement of constraining edges on triangulations of smooth surfaces. We prove several combinatorial properties of restricted CDTs, including conditions under which the restricted CDT contains every constraining segment, conditions under which the restricted CDT is homeomorphic to the underlying surface Σ, and a characterization of which vertices must be considered to compute the triangles near a segment. The restricted CDT has immediate practical applications in surface meshing and geometric modeling. Along the way we improve many commonly used supporting results in geometric sampling theory.

In the second part of this thesis we apply the geometric tools and high-dimensional intuition developed in the previous chapters to problems in machine learning. We study the problem of adversarial examples, a pervasive phenomenon of machine learning models where perturbations of the input that are imperceptible to humans reliably lead to confident incorrect classifications. We study robustness to adversarial examples under the “Manifold Hypothesis”: the observation that ‘real’ data often exhibits low-dimensional structure. Our results highlight the role of codimension, the difference between the dimension of the data manifold and the dimension of the embedding space, in adversarial robustness. We prove a tradeoff between robustness in different norms, show that adversarial training is sample inefficient, and that robustness requires larger models.

Lastly we study the relationship between robustness and optimization in the linear regression setting. We show an example of a learning problem for which the solution found by adaptive optimization algorithms exhibits qualitatively worse robustness properties against both L_2- and L_∞-adversaries

than the solution found by non-adaptive algorithms. Then we fully characterize the geometry of the loss landscape of L_2-adversarial training in least-squares linear regression. The geometry of the loss landscape is subtle and has important consequences for optimization algorithms.

Dedicated to my parents, Antoinette and Isaac Khoury, and younger brother, Jason Khoury.

Contents

List of Figures

Acknowledgments

First and foremost, I would like to thank my advisor, Jonathan Shewchuk. Surely this thesis would not have been possible without many of your brilliant ideas. Thank you for showing me how you think about problems. Many of the techniques I learned from you have become indispensable tools. Thank you for your incredible patience and for providing me the freedom to explore my own ideas. More importantly, thank you for being my friend. Our conversations over the years are some of my fondest memories in graduate school.

I would like to thank Vladlen Koltun for a wonderful year as an intern in the Intelligent Systems Lab at Intel Research. You taught me the basics of experimental work and pushed me to become a better researcher. For that I am grateful. My work in your group was my first introduction to deep learning, which would go on to be an important topic in this dissertation. My time at Intel was also greatly influenced by my collaborator Qian-Yi Zhou. Thank you for providing much needed encouragement and ideas at critical points throughout my internship. Lastly I'd like to thank the rest of the group at the time including René Ranftl, Jaesik Park, Jia Xu, Vibhav Vineet, Qifeng Chen, Alexey Dosovitskiy, Arno Knapitsch, and my office mate Sohil Shah.

I would like to thank my undergraduate advisor, Rephael Wenger. Thank you for showing me the beauty of geometric algorithms and sparking my interest in research. Like Jonathan, you were always remarkably patient and encouraging. Surely I would not have been admitted to the PhD program at UC Berkeley had it not been for our fateful encounter so many years ago.

To Dylan Hadfield-Menell, I remember the exact moment that I knew you and I would become friends. Truly getting there took a bit longer than I would have liked, but it was worth it. You have been my best friend throughout graduate school. Thank you for the bike rides, the sour beers, and for teaching me how to ski. Thank you for pointing me in the direction of adversarial robustness and for your guidance on our first project together. More importantly, thank you for being my confidant during the more pressing times of graduate school. I hope I was able to provide the same to you.

To Judy Savitskaya, there simply are no words. Nothing I write could possibly do justice to our friendship over the last eight years. I love you, unconditionally.

My time at UC Berkeley was made so much more enjoyable by the many amazing individuals that I befriended along the way. Thank you to Horia Mania for the drinks. Tess Smidt for the friendsgivings. Sean Lubner for teaching me how to deadlift safely. Melih Elibol for the steak. Tom, Rhonda and Emily Davenport for welcoming me into your home, the sense of family, and the bike rides. Robert Nishihara for the interesting questions that made me think. Matthew Mirman for your unpredictability. Jonathan Lin for the many enjoyable hours we spent together in Asha Tea House. Smitha Milli for the comedy, the tea, and the muffins. Katelin Schutz for the cake. Sandy Huang for the cupcakes. (It's important to have people in your life that know how to make quality baked goods.) Michael Schoenberg, Chris Mogni, and Marie Hepfer for just being there when I needed you.

Finally, I would like to thank my parents, Antoinette and Isaac Khoury, as well as my brother, Jason Khoury, for their support. Thank you for encouraging me to pursue my dreams and doing everything you could to help me along the way. No matter what situation I find myself in, I know I can always count on you to help me through it. I dedicate this thesis to you.

Chapter 1

Introduction

In the first part of this thesis we study a fundamental problem in computational geometry. Manifold triangulation is the problem of, given a set of points V sampled densely from some known manifold Σ, output a set of simplices $\mathcal{T}$ that is topologically identical to Σ and geometrically close to Σ. By "topologically identical" we mean that there exists a homeomorphism h (a continuous bijection with continuous inverse) between $\mathcal{T}$ and Σ. If such a homeomorphism exists, $\mathcal{T}$ and Σ have the "same shape" up to topology. However just because two spaces are topologically identical does not mean that they have similar geometry. (A tea cup is topologically identical to a coffee cup but they look very different.) By "geometrically close" we mean that $\mathcal{T}$ is everywhere close to Σ in Euclidean distance and the normal spaces of $\mathcal{T}$ well approximate those of Σ. The first requirement can be expressed by a condition on the homeomorphism h that $\|h(x) - x\| \leq \epsilon$ for all $x \in \mathcal{T}, h(x) \in \Sigma$; the second by a similar condition on the normal spaces.

In Chapter 3 we study a variant of this problem with one additional constraint. In addition to a set of points V densely sampled from a known Σ we are also given a finite set of line segments S, whose endpoints are in V, which must appear as edges in the output triangulation $\mathcal{T}$. We study this problem in the case where Σ is a surface (2-manifold) embedded in $\mathbb{R}^3$. To solve this problem we introduce *restricted constrained Delaunay triangulations* (restricted CDTs), which combine ideas from constrained Delaunay triangulations and restricted Delaunay triangulations to enable the enforcement of constraining edges on triangulations of smooth surfaces. We prove several combinatorial properties of restricted CDTs, including conditions under which the restricted CDT contains every constraining segment, conditions under which the restricted CDT is homeomorphic to the underlying surface Σ, and a characterization of which vertices must be considered to compute the triangles near a segment. The restricted CDT has immediate practical applications in surface meshing and geometric modeling.

To prove many of the properties of restricted CDTs, we needed to improve many commonly used supporting results in geometric sampling theory. In Chapter 2 we present several results which can be used to bound the *normal error* between $\mathcal{T}$ and Σ. Consider a triangle $\tau \in \mathcal{T}$ whose vertices are in V. Let $x \in \tau$ be any point on the triangle, not necessarily a vertex, and let $\hat{x} = h(x)$. Then the error between the normal vector of τ and the normal vector to Σ at $\hat{x}$ can be decomposed as $\angle(n_\tau, n_{\hat{x}}) \leq \angle(n_\tau, n_v) + \angle(n_v, n_{\hat{x}})$ where v is the vertex of τ with largest plane angle. (Note that all three angles are measured as the acute angle between the relevant vectors.) The first term, $\angle(n_\tau, n_v)$, is called the triangle normal error, the error of the normal vector of the triangle to the normal vector to Σ at v. We prove a sharp bound on $\angle(n_\tau, n_v)$ in terms of the size and shape of τ. Our proof is very intuitive and is sharp, meaning that the bound is tight including constants and an example exists that attains the bound. The second term, $\angle(n_v, n_{\hat{x}})$ is called the normal variation, the error incurred as the normal varies along the surface of Σ. The rate at which the normal varies is dependent upon the curvature of Σ. When v and $\hat{x}$ are close in Euclidean distance, relative to the local curvature, we can bound the normal variation. Unlike our Triangle Normal Lemma which only applies to triangles, our Normal Variation Lemmas apply to k-dimensional smooth manifolds embedded $\mathbb{R}^d$, for any choice of $k < d$.

In the second part of this thesis we apply the geometric tools and high-dimensional intuition developed in the previous chapters to problems in machine learning. We study the problem of adversarial examples, a pervasive phenomenon of machine learning models where perturbations of the input that are imperceptible to humans reliably lead to confident incorrect classifications. Given a labeled training set, the goal of classification is to output a classifier f that performs well on the underlying distribution from which the training set was sampled. The existence of adversarial examples implies that there exists a small perturbation δ such that $f(x) \neq f(x + \delta)$. Geometrically, starting at a correctly classified point, there is a direction in which if we walk a small distance we cross the decision boundary.

The classifier f induces a decision boundary in the input space. A classifier is robust to adversarial examples if the decision boundary it induces is as far from the data distribution as possible, with respect to a relevant distance metric. In Chapter 4 we study robustness to adversarial examples under the "Manifold Hypothesis": the observation that 'real' data often exhibits low-dimensional structure. We model data as being sampled from class-specific low-dimensional manifolds embedded in a high-dimensional space. We consider a threat model wherein an adversary may choose any point on the data manifold to perturb by ϵ in order to fool a classifier. To be robust to such an adversary, a classifier must be correct everywhere in an ϵ-tube around the data manifold. We highlight the role of codimension, the difference between the dimension of the data manifold and the dimension of the embedding space, in adversarial robustness. Furthermore we prove a tradeoff between robustness in different norms, show that adversarial training is sample inefficient, and that robustness requires larger models. We also present adversarial training with Voronoi constraints, a modification to the standard adversarial training paradigm which we show improves robustness in high-codimension settings.

In Chapter 5 we study the relationship between robustness and optimization in the linear regression setting. We show an example of a learning problem for which the solution found by adaptive optimization algorithms exhibits qualitatively worse robustness properties against both L_2- and L_∞-adversaries than the solution found by non-adaptive algorithms. The robustness of the adaptive solution decreases rapidly as the dimension of the problem increases, while the robustness of the non-adaptive solution is stable as the dimension increases. Then we fully characterize the geometry of the loss landscape of L_2-adversarial training in least-squares linear regression. The L_2-adversarial training objective is convex everywhere; moreover, it is strictly convex everywhere except along either 0, 1, or 2 line segments, depending on the value of ϵ. Furthermore for nearly all choices of ϵ, these line segments along which the objective is convex, but not strictly convex, lie outside of the rowspace and the gradient along these line segments is nonzero. Surprisingly the solution is almost always unique, and thus common optimization algorithms are guaranteed to converge to the robust solution.

Chapter 2

Approximation Bounds for Normals on Triangulated Surfaces and Manifolds

2.1 Introduction

Triangulations of surfaces are used heavily in computer graphics, visualization, and geometric modeling; they also find applications in scientific computing. Also useful are triangulations of manifolds in spaces of dimension higher than three—for example, as a tool for studying the topology of algebraic varieties. A surface triangulation (sometimes called a *surface mesh*) replaces a curved surface with flat triangles—or in higher dimensions, simplices—which are easy to process and suitable for graphics rendering engines; but they introduce error. How good is a triangulation as an approximation of a curved surface?

The two criteria most important in practice are the *interpolation error*, the error in the position of the surface, and the *normal error*, the error in the normal vectors of the surface. Let Σ be a surface or manifold embedded in a Euclidean space $\mathbb{R}^d$, and let Λ be a piecewise linear surface or manifold formed by a triangulation that approximates Σ. The interpolation error can be quantified as the distance from an arbitrary point on Λ to the nearest point on Σ, or vice versa. The normal error can be quantified by choosing two nearby points $x \in \Lambda$ and $y \in \Sigma$—a natural choice of y is the point on Σ nearest x—and measuring the angle separating the vector normal to Λ at x from the vector normal to Σ at y. (The vector normal to Λ is usually undefined if x lies on a boundary where simplices meet, but our results will treat simplices individually rather than treat Λ as a whole.)

Some notation: we employ a correspondence between the two surfaces called the *nearest-point map*[1] ν, which maps a point $x \in \mathbb{R}^d$ to the point $\nu(x)$ nearest x on Σ (if that point is unique). We will frequently use the abbreviation $\tilde{x}$ to denote $\nu(x)$. Given two points $p, q \in \mathbb{R}^d$, pq denotes a line segment with endpoints p and q, and $|pq|$ denotes its Euclidean length $\|p - q\|_2$. For a point p on a surface $\Sigma \subset \mathbb{R}^3$, n_p denotes a vector normal to Σ at p (whose magnitude is irrelevant). For a triangle $\tau \subset \mathbb{R}^3$, n_τ denotes a vector normal to τ. Let $\angle(n_\tau, n_p)$ denote the angle separating n_τ from n_p. In

[1] We follow the convention of Cheng et al. [20] and use the Greek letter *nu*, which unfortunately is hard to distinguish from the italic Roman letter v.

higher-dimensional Euclidean spaces, the normal vectors may be replaced by normal subspaces; see Section 2.2.

The goal of this chapter is to provide strong bounds on the normal errors for triangles, based on assumptions about the sizes of medial balls (defined in Section 2.2). Specifically, given a triangle τ whose vertices lie on Σ and a point $x \in \tau$, we bound the angle $\angle(n_\tau, n_{\tilde{x}})$. Besides the normal errors, we also study the *normal variation*, the angle separating the normal vectors (or normal spaces) at two different points on Σ. (We need to understand the normal variation to study the normal error; it is also used to prove that certain triangulations are homeomorphic to a surface [20, 28].) Bounds on both of these quantities— the normal error and the normal variation—have been derived in prior works [2, 6, 5, 19, 20, 28] and form a foundation for the correctness and accuracy of many algorithms in surface reconstruction [2, 6, 4, 5, 10, 19, 31, 28, 58] and mesh generation [13, 11, 21, 20, 32, 69, 74] based on Delaunay triangulations. Our improved bounds directly imply improved *sampling bounds* for all of those algorithms. By "sampling bounds," we mean how densely points must be sampled on a surface to guarantee that the reconstructed surface or the surface mesh has a good approximation accuracy and the correct topology.

A second goal of this chapter is to generalize our bounds to manifolds in higher dimensions. Our bounds on the normal error apply only to triangles, albeit on a manifold of any dimension (greater than 1) in a space of any dimension. (We would like to study normal errors for simplices of higher dimension, but the interaction between the shape of, say, a tetrahedron in $\mathbb{R}^4$ and the stability of its normal space is complicated. It deserves more study.)

Our bounds on the normal variation also apply in higher dimensions, but with a twist. The *codimension* of a k-manifold $\Sigma \subset \mathbb{R}^d$ is $d - k$. We have two *normal variation lemmas* (Section 2.4): one for codimension 1, which bounds an angle $\angle(n_p, n_q) \in [0°, 180°]$ between two normal vectors, and one for higher codimensions, which bounds an angle $\angle(N_p\Sigma, N_q\Sigma) \in [0°, 90°]$ between two normal spaces (see Section 2.2 for definitions of normal spaces and the angles between them). The reason for two separate lemmas is that the codimension 1 bound is stronger; codimension 2 introduces configurations that weaken the bound and cannot occur in codimension 1. As a consequence, some of our bounds on the normal errors also depend on the codimension.

One of our results on the normal error improves a prior bound by a factor of about 1.9 (see Section 2.3). Even small constant-factor improvements in the bounds are valuable; for example, the number of triangles necessary for a surface mesh to guarantee a specified accuracy in the normals is reduced by a factor of $1.9^2 = 3.61$, helping to substantially speed up the application using the mesh. In dimensions higher than three, we are not aware of prior bounds with explicitly stated constants, but there are asymptotic results [19]; part of our contributions is to give strong explicit bounds. Our bound on the interpolation error is *sharp*, meaning that it cannot be improved (without making additional assumptions). We conjecture that our bound on the normal variation in codimension 1 is *sharp*, meaning that it cannot be improved (without making additional assumptions). (We use *sharp* to mean that not even the constants can be improved, as opposed to *tight*, which is sometimes used in an asymptotic sense.)

The bounds help to clarify the relationship between approximation accuracy, the sizes and shapes of the simplices in a surface mesh, and the geometry of the surface itself. Reducing the sizes of the simplices tends to reduce both the interpolation and normal errors; unsurprisingly, finer meshes offer better approximations than coarser ones. The interpolation errors on a simplex scale quadratically with the size of the simplex. This is good news: shrinking the simplices reduces the interpolation error quickly. The normal errors scale linearly (not quadratically) with the size of the simplex. Roughly speaking, both types of error scale linearly with the curvature of the manifold, measured at a selected point; more precisely, they scale inversely with the radii of selected medial balls (defined in Section 2.2), which we use to impose appropriate bounds on both the curvature and the proximity of different parts of a manifold. Therefore, portions of a manifold with greater curvature require smaller simplices.

Normal errors are very sensitive to the shape of a simplex. Skinny simplices underperform simplices that are close to equilateral, and really skinny simplices can yield catastrophically wrong normals. As a rough approximation, the worst-case normal error on a triangle is linearly proportional to the triangle's circumradius, defined in Section 2.2. (See Sections 2.3 and 2.5 and Amenta, Choi, Dey, and Leekha [6]). For triangles with a fixed longest edge length, the worst normal errors are suffered by triangles with angles close to 180°, because the circumradius approaches infinity as the largest angle approaches 180°. We give several bounds on the normal error for a triangle: the simplest one depends on the triangle's circumradius, whereas a stronger bound depends on one of triangle's angles as well, giving us a more nuanced understanding of the relationship between triangle shape and normal errors.

2.2 A Tour of the Bounds

To create a surface mesh that meets specified constraints on accuracy, one must consider the geometry of Σ and the size and (sometimes) the shape of each simplex. Our bounds use three parameters to measure a simplex τ: the min-containment radius of τ and, for triangles only, the circumradius of τ and (optionally) one of τ's plane angles.

For a simplex $\tau \subset \mathbb{R}^d$, the *smallest enclosing ball* of τ (also known as the *min-containment ball*) is the smallest closed d-dimensional ball $B_\tau \supseteq \tau$, illustrated in Figure 2.1. The *min-containment radius* of τ is the radius of τ's smallest enclosing ball; we write it as r (though sometimes r will be the radius of any arbitrary enclosing ball). The *diametric ball* of τ is the smallest closed d-ball B such that all τ's vertices lie on B's boundary, also illustrated in Figure 2.1. The *circumcenter* and *circumradius* of τ are the center and radius of τ's diametric ball, respectively; we write the circumradius as R. For every simplex, $r \leq R$; but if τ is "badly" shaped, R can be arbitrary large compared to r. (Recall that for a triangle, $R \to \infty$ as the largest angle approaches 180° and the longest edge remains fixed.) A simplex τ always contains the center of its smallest enclosing ball, but frequently not its circumcenter. The center of τ's smallest enclosing ball is the point on τ closest to τ's circumcenter. (See Rajan [73, Lemma 3] for an algebraic proof based on quadratic program duality, or Shewchuk [82, Lemma 24] for a geometric proof.) Hence, $r = R$ if and only if τ contains its circumcenter.

The *circumcircle* (circumscribing circle) of a triangle $\tau \subset \mathbb{R}^d$ is the unique circle that passes through all three vertices of τ. The circumcircle has the same center and radius R as τ's diametric ball (i.e., τ's circumcenter and circumradius). A *plane angle* of a triangle τ is one of the usual three

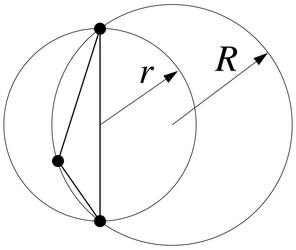

Figure 2.1: The smallest enclosing ball of a triangle, with radius r, and the triangle's diametric ball, with radius R.

angles we associate with a triangle, though τ might be embedded in a high-dimensional space. A triangle contains its circumcenter (and has $r = R$) if and only if it has no plane angle greater than 90°.

There are two salient aspects to the geometry of Σ. One is curvature: a surface with greater curvature needs smaller triangles. (Nonsmooth phenomena like sharp edges can make the triangle normals inaccurate no matter how small the triangles are, and are best addressed by matching the triangle edges to the surface discontinuities. We don't address that problem here.) A more subtle aspect is that a surface can "double back" and come close to itself in Euclidean space: for example, if a mesh of a hand has a triangle connecting the pad of the thumb to a knuckle of the index finger, the triangle misrepresents the surface badly.

The early literature on provably good surface reconstruction identified the *medial axis*—more specifically, the sizes of medial balls—as an effective way to gauge the triangle sizes required as a consequence of both curvature and the proximity of parts like fingers. Let Σ be a bounded, smooth k-manifold embedded in $\mathbb{R}^d$. Let $B \subset \mathbb{R}^d$ be an open ball. We call B *surface-free* if $B \cap \Sigma = \emptyset$. We say B *touches* Σ if $B \cap \Sigma = \emptyset$ but B's boundary intersects Σ; that is, B is surface-free but its closure is not. In that case, B is tangent to Σ at the intersection point(s). There are two types of *medial ball*; both types are surface-free balls that touch Σ, as illustrated in Figure 2.2. Every surface-free ball whose boundary touches Σ at more than one point is a medial ball; most medial balls are of this first type. Let $W \subset \mathbb{R}^d$ be the set containing the center of every medial ball of this first type; these are the points $w \in W$ where the nearest-point map $\nu(w)$ is not uniquely defined. (Recall that ν maps a point $x \in \mathbb{R}^3$ to the point $\tilde{x} = \nu(x)$ nearest x on Σ.) The *medial axis* $M \in \mathbb{R}^d$ is the closure of W, as illustrated. Each point added to M by taking the closure is the center of a medial ball of the second type, which touches Σ at just one point.

We will often refer to the medial balls tangent to Σ at a point $p \in \Sigma$. In codimension 1, there are typically two such balls (but sometimes just one), one enclosed by Σ and (optionally) one outside Σ. In higher codimensions, there are infinitely many. All their centers lie in the normal space $N_p\Sigma$. A useful construction we will use later is to choose a point $q \in N_p\Sigma \setminus \{p\}$ and imagine an open ball tangent to Σ at p whose radius is initially zero; then the ball grows so that its center moves along the ray $\vec{pq}$ while its boundary remains touching p. Typically, at some point the ball will not be able to grow further without intersecting Σ. At the last instant when the ball is still surface-free, it is a

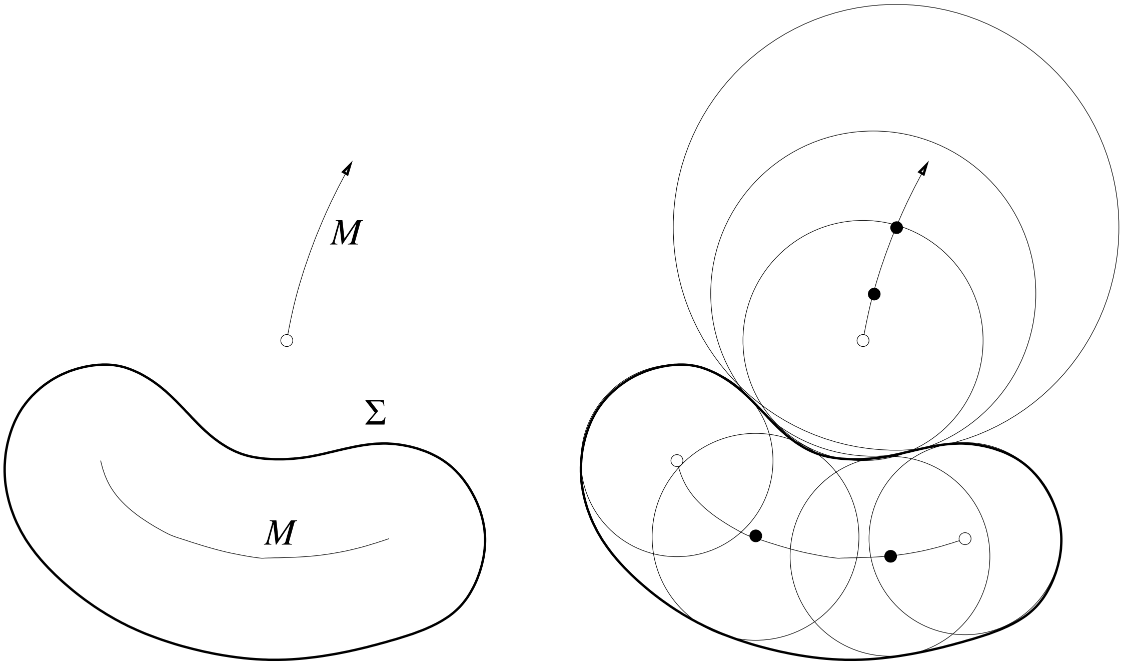

Figure 2.2: Left: A 1-manifold Σ and its medial axis M. Right: Some of the medial balls that define M. Those with black centers are medial balls of the first type; those with white centers are of the second type.

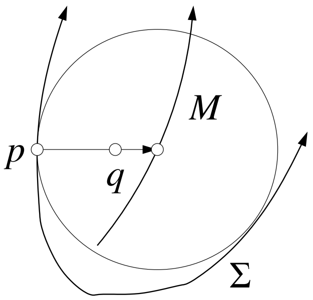

Figure 2.3: The medial ball tangent to Σ at p whose center lies on the ray $\vec{pq}$.

medial ball, and its center is a point in the medial axis M. Typically the ball cannot grow further because it touches a second point on Σ (producing a medial ball of the first type), but sometimes it is constrained solely by the curvature of Σ at p itself (producing a medial ball of the second type). In some cases when p lies on the boundary of the convex hull of Σ, the ball can grow to infinite radius and degenerate into an open halfspace while remaining surface-free. It is occasionally useful to refer to such a degenerate medial ball, although it does not contribute a point to M.

For any $p \in \Sigma$, the *empty ball size* ebs(p) is the radius of the smallest medial ball tangent to Σ at p. The *local feature size* lfs(p) is the distance from p to the medial axis (i.e., from p to the nearest point on M). Formally,

$$\mathrm{lfs}(p) = \min_{m \in M} |pm|; \qquad \mathrm{ebs}(p) = \min_{m \in M \cap N_p\Sigma} |pm|.$$

This definition makes clear that $\mathrm{lfs}(p) \leq \mathrm{ebs}(p)$. Both measures simultaneously constrain the curvature of Σ at p (the principle curvatures cannot exceed $1/\mathrm{ebs}(p)$) and the proximity of other "parts" of the manifold (recall the example of fingers of a hand). The empty ball size has the advantage that it is more local in nature than the local feature size, so bounds expressed in terms of $\mathrm{ebs}(p)$ are more generally applicable (which is why we are introducing ebs here). The local feature size $\mathrm{lfs}(p)$ constrains the curvature not only at p, but also at nearby points, permitting the proof of stronger conclusions. The local feature size is 1-Lipschitz, meaning that for all $p, q \in \Sigma$, $\mathrm{lfs}(p) \leq \mathrm{lfs}(q) + |pq|$; whereas the empty ball size can vary rapidly over Σ.

One of the main contribution of the early literature on provably good surface reconstruction was to recognize that the local feature size (scaled down by a constant factor) is a good guide to how closely points need to be spaced on Σ to ensure that surface reconstruction algorithms will produce a correct output that approximates Σ well [2, 3]. Subsequently, provably good surface mesh generation algorithms also adopted these observations [12, 11, 18].

The normal errors are (approximately) inversely proportional to $\mathrm{ebs}(p)$ or $\mathrm{lfs}(p)$ for some relevant point p. That is, the errors increase with a decreasing radius of curvature (i.e., an increasing curvature). If Σ is not smooth, each point p where Σ is not smooth has $\mathrm{ebs}(p) = \mathrm{lfs}(p) = 0$, and p lies on the medial axis M. Our bounds do not apply at such points (the bounds are infinite). The bounds still apply at points where ebs is positive.

Before we discuss normal errors, we must discuss our Normal Variation Lemmas (Section 2.4). The smoothness of a manifold Σ implies that if two points are close to each other, their normal spaces differ by only a small angle, and likewise for their tangent spaces. Given two points $p, q \in \Sigma$, a *normal variation lemma* gives an upper bound on the angle between their normal vectors (in codimension 1) or their normal spaces (in codimension 2 or higher).

What are tangent spaces and normal spaces? A *k-flat*, also known as an *k-dimensional affine subspace*, is a k-dimensional space that is a subset of $\mathbb{R}^d$. It is essentially the same as a k-dimensional subspace (from linear algebra), but whereas a subspace must contain the origin, a flat has no such requirement. Given a smooth k-manifold $\Sigma \subset \mathbb{R}^d$ and a point $p \in \Sigma$, the *tangent space* $T_p\Sigma$ is the k-flat tangent to Σ at p, and the *normal space* $N_p\Sigma$ is the $(d-k)$-flat through p that is entirely orthogonal (complementary) to $T_p\Sigma$; that is, every line in $N_p\Sigma$ is perpendicular to every line in $T_p\Sigma$.

Recall that the codimension of Σ is $d - k$. In the special (but common) case of codimension 1, a $(d-1)$-manifold without boundary divides $\mathbb{R}^d$ into an unbounded region we call "outside" and one or more bounded regions we call "inside." Hence for codimension 1 we use the convention that any normal vector n_p is directed outward. The normal space $N_p\Sigma$ is a line parallel to n_p, but n_p is directed and $N_p\Sigma$ is not. In codimension 2 or higher, the normal space has dimension 2 or higher (matching the codimension of Σ) and Σ might not even be orientable, so we don't assign $N_p\Sigma$ a direction.

Let $F, G \subseteq \mathbb{R}^d$ be two flats, and suppose that the dimension of F is less than or equal to the dimension of G. We define the angle separating F from G to be

$$\angle(F, G) = \angle(G, F) = \max_{\ell_F \subset F} \min_{\ell_G \subset G} \angle(\ell_F, \ell_G)$$

where ℓ_F and ℓ_G are lines. Note that if F and G are of different dimensions, the "max" must apply over the lower-dimensional flat and the "min" over the higher-dimensional flat. This angle is always in the range $[0°, 90°]$; we use angles greater than 90° only for directed vectors. If $F_\perp$ denotes a flat complementary to F, it is well known that $\angle(F, G) = \angle(G_\perp, F_\perp)$; hence, for two points $p, q \in \Sigma$, $\angle(N_p\Sigma, N_q\Sigma) = \angle(T_p\Sigma, T_q\Sigma)$. Note that there is more than one way to define "angles between subspaces." The best-known way originates with an 1875 paper of Jordan [52]; by this reckoning, one needs multiple angles to fully characterize the angular relationships between two high-dimensional flats. Our definition corresponds to the greatest of these angles (including the 90° angles, which are not included in Jordan's *canonical angles*), so our upper bound holds for all the angles.

It is convenient to specify our bounds on $\angle(N_p\Sigma, N_q\Sigma) = \angle(T_p\Sigma, T_q\Sigma)$ in terms of a parameter $\delta = |pq|/\mathrm{lfs}(p)$. The worst-case value of $\angle(N_p\Sigma, N_q\Sigma)$ is $\delta + O(\delta^3)$ radians for small δ. Hence, the worst-case normal variation is approximately linear in $|pq|$ and approximately inversely proportional to $\mathrm{lfs}(p)$.

We give two Normal Variation Lemmas that, collectively, apply to smooth k-manifolds embedded in $\mathbb{R}^d$ for every d and $k < d$. They are stronger than the best prior bounds, especially for $d > 3$. There are two separate lemmas because we obtain a better bound for codimension 1 than for codimension 2 and higher. Our main result in codimension 1 is that for $\delta \leq 0.9717$, $\angle(n_p, n_q) \leq \eta_1(\delta) \in [0°, 180°]$ where

$$\eta_1(\delta) = \arccos\left(1 - \frac{\delta^2}{2\sqrt{1-\delta^2}}\right) \approx \delta + \frac{7}{24}\delta^3 + \frac{123}{640}\delta^5 + \frac{1{,}083}{7{,}168}\delta^7 + O(\delta^9).$$

Our main result for general codimensions is that for $\delta \leq 0.7861$, $\angle(N_p\Sigma, N_q\Sigma) = \angle(T_p\Sigma, T_q\Sigma) \leq \eta_2(\delta) \in [0°, 90°]$ where

$$\eta_2(\delta) = \arccos\sqrt{1 - \frac{\delta^2}{\sqrt{1-\delta^2}}} \approx \delta + \frac{5}{12}\delta^3 + \frac{57}{160}\delta^5 + \frac{327}{896}\delta^7 + O(\delta^9).$$

We conjecture that our bound for codimension 1 is sharp, meaning that it cannot be improved without imposing additional restrictions. Our bound for codimension 2 is not sharp and leaves room for improvement. See Section 2.4 for additional bounds (and plots thereof) that are stronger when the distance from q to p's tangent plane is known.

Figure 2.4 compares our two bounds and two prior bounds for surfaces in $\mathbb{R}^3$, both by Amenta and Dey [5]. The stronger prior bound is $\angle(n_p, n_q) \leq -\ln(1-\delta)$ radians for $\delta \leq 0.9567$. (A derivation of both bounds can also be found in Cheng et al. [20]. Amenta and Bern [2] gave an early normal variation lemma with a weaker bound, but the proof was erroneous.) This bound fades to 90° at $\delta \approx 0.7921$ and to 180° at $\delta \approx 0.9567$, whereas our bound for codimension 1 fades to 90° at $\delta \approx 0.9101$ and to 180° at $\delta \approx 0.9717$. Our bound for higher codimensions fades to 90° at $\delta \approx 0.7861$ and stops there (because we do not assign directions to normal spaces of dimension 2 or higher).

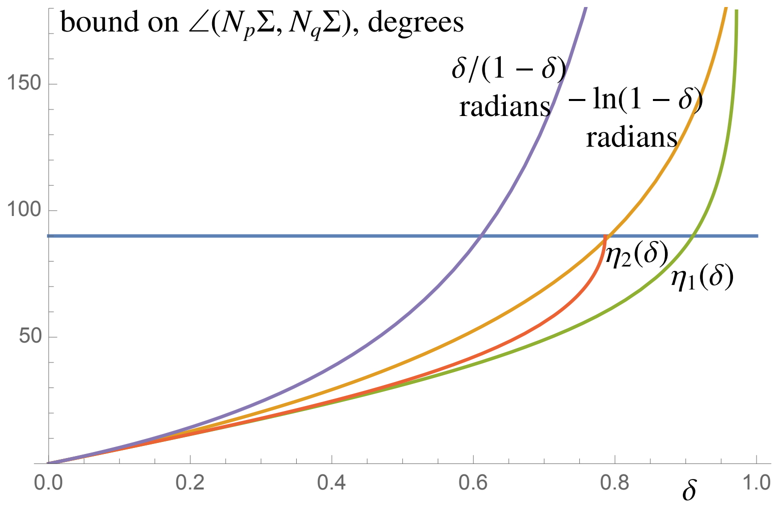

Figure 2.4: Upper bounds in degrees for $\angle(N_p\Sigma, N_q\Sigma)$ as a function of $\delta = |pq|/\mathrm{lfs}(p)$, provided by several normal variation lemmas. The brown curve is the bound $-\ln(1-\delta)$ radians proved by Amenta and Dey [5] for surfaces without boundary in $\mathbb{R}^3$. The purple curve is the weaker but better-known bound $\delta/(1-\delta)$ radians, also by Amenta and Dey [5]. The green curve is our bound for codimension 1—that is, for $(d-1)$-manifolds without boundary in $\mathbb{R}^d$. The red curve is our bound for codimension 2 or greater—that is, for k-manifolds without boundary in $\mathbb{R}^d$ with $d-k \geq 2$. Bounds between 90° and 180° are meaningful for manifolds without boundary in codimension 1. The red curve stops at 90° because we do not assign directions to normal spaces of dimension 2 or higher.

Amenta and Dey [5] also proved a bound of $\delta/(1-\delta)$ radians, which has become better known. We include it in Figure 2.4 (in purple) to show how much is lost by using the well-known bound instead of the stronger bounds. The Amenta–Dey bounds are of the form $\angle(N_p\Sigma, N_q\Sigma) \leq \delta + O(\delta^2)$ radians, whereas our bounds show that $\angle(N_p\Sigma, N_q\Sigma) \leq \delta + O(\delta^3)$ radians.

Cheng, Dey, and Ramos [19] prove a general-dimensional normal variation lemma for k-manifolds in $\mathbb{R}^d$, showing that in the worse case, $\angle(N_p\Sigma, N_q\Sigma)$ grows linearly with δ for small δ; but they express their bound in an asymptotic form with an unspecified constant coefficient, which makes a comparison with our bounds difficult. We think it is a useful and practical contribution to provide explicit numerical bounds $\eta_1(\delta)$ and $\eta_2(\delta)$ for $d > 3$. Although our bound $\eta_2(\delta)$ is not sharp, for $\delta \leq 0.7$ it is not much bigger than $\eta_1(\delta)$, which we conjecture is a lower bound for all codimensions.

Finally, our results include several Triangle Normal Lemmas (Sections 2.3 and 2.5). For a triangle τ whose vertices lie on a k-manifold Σ, let $\nu(\tau)$ be the image of τ under the nearest-point map. We derive bounds on how well τ's normal vector locally approximates the vectors normal to Σ on $\nu(\tau)$. For a j-simplex $\tau \subset \mathbb{R}^d$, τ's tangent space is its affine hull, a j-flat denoted aff τ. For convenience, we define a particular normal space for simplices: let N_τ denote the set of points in $\mathbb{R}^d$ that are equidistant to all the vertices of τ. N_τ is a $(d-j)$-flat complementary to aff τ. The intersection of N_τ and aff τ is τ's circumcenter.

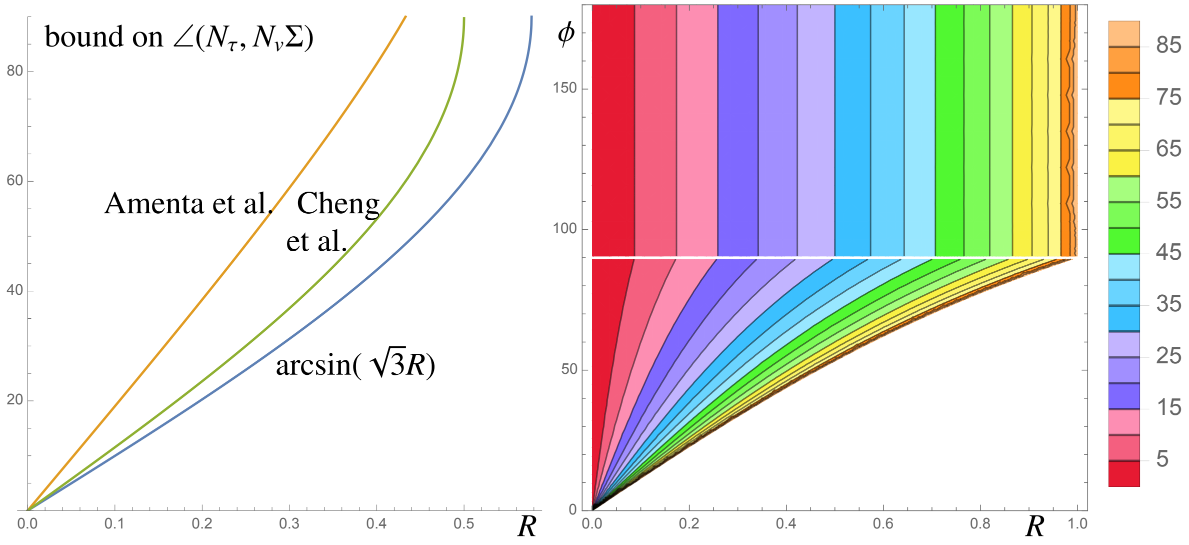

Figure 2.5: Upper bounds in degrees for $\angle(N_\tau, N_v\Sigma) = \angle(\operatorname{aff}\tau, T_v\Sigma)$, where τ is a triangle whose vertices lie on a manifold Σ and v is a vertex of τ. We assume $\operatorname{ebs}(v) = 1$. Left: three bounds on $\angle(N_\tau, N_v\Sigma)$ for the case where v is the vertex at τ's largest plane angle (or any angle 60° or greater), as a function of the circumradius R of τ. The blue curve is our new bound (2.2). The green curve is the best (albeit little-known) prior bound we are aware of, $\arcsin(2R)$, due to Cheng, Dey, Edelsbrunner, and Sullivan [18]. The brown curve is a much better-known prior bound, due to Amenta, Choi, Dey, and Leekha [6] (see Lemma 2). Right: isocontour plot of our bound (2.1) as a function of the circumradius R (on the horizontal axis) and the angle ϕ at the vertex v (on the vertical axis). For small ϕ, the lemma does not provide a bound (unless R is very small), but see Section 2.5.

Our basic Triangle Normal Lemma applies only at the vertices of τ. Let R be τ's circumradius. Let v be a vertex of τ and let ϕ be τ's plane angle at v. Then

$$\angle(N_\tau, N_v\Sigma) = \angle(\operatorname{aff}\tau, T_v\Sigma) \le \arcsin\left(\frac{R}{\operatorname{ebs}(v)} \max\left\{\cot\frac{\phi}{2}, 1\right\}\right). \tag{2.1}$$

Note that the argument $\cot\frac{\phi}{2}$ dominates if ϕ is acute and the argument 1 dominates if ϕ is obtuse. If v is the vertex at τ's largest plane angle (so $\phi \ge 60°$), then

$$\angle(N_\tau, N_v\Sigma) = \angle(\operatorname{aff}\tau, T_v\Sigma) \le \arcsin\frac{\sqrt{3}R}{\operatorname{ebs}(v)}. \tag{2.2}$$

Figure 2.5 plots both bounds, (2.2) at left and (2.1) at right. Note that $\operatorname{ebs}(v)$ can be replaced by $\operatorname{lfs}(v)$. It is interesting that the worst case preventing the bound (2.2) from being better is incurred by an equilateral triangle (rather than a triangle with a very large or small angle, as one might expect).

These bounds vary approximately linearly with the circumradius of τ, and inversely with the empty ball size or local feature size at v. Whereas the interpolation error varies quadratically with the radius of τ's smallest enclosing ball, and is therefore very sensitive to τ's size but nearly

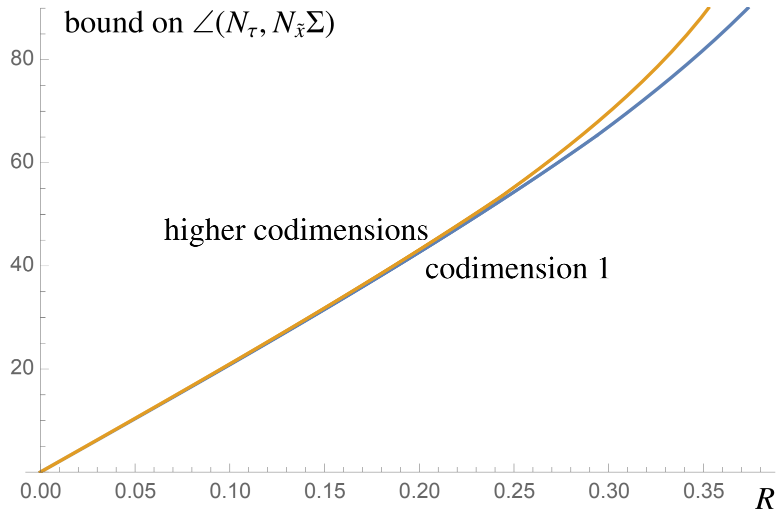

Figure 2.6: Upper bounds for $\angle(N_\tau, N_{\tilde{x}}\Sigma) = \angle(\text{aff}\,\tau, T_{\tilde{x}}\Sigma)$ as a function of the circumradius R of τ, where τ is a triangle whose vertices lie on a manifold Σ and x is any point on τ. We assume all three vertices w of τ satisfy $\text{lfs}(w) \geq 1$. The blue curve is the upper bound in codimension 1 (with the choice $\phi = 49°$) and the brown curve is the upper bound in higher codimensions (with the choice $\phi = 48.5°$), for which the Normal Variation Lemma is weaker.

insensitive to its shape, the normal error varies (linearly) with τ's circumradius, which can be much larger than τ if τ has a large angle (close to 180°). It is well known that in surface meshes, triangles with large angles are undesirable and sometimes even crippling to applications, not because of problems with interpolation error, but because of problems with very inaccurate normals.

Given a triangulation of Σ, one would like to have a triangle normal lemma that applies to every point on Σ, not just at the vertices. Moreover, the Triangle Normal Lemma bounds are weak or nonexistent at the vertices where the triangles have small plane angles. Hence, we use the Normal Variation Lemmas to extend the Triangle Normal Lemma bounds over the rest of $\nu(\tau)$—that is, for every $x \in \tau$, we bound $\angle(N_\tau, N_{\tilde{x}}\Sigma)$. Thus, a finely triangulated smooth manifold accurately approximates the normal spaces of all the points on the manifold. We call these results *extended triangle normal lemmas*. Suppose that $R \leq \kappa\,\text{lfs}(w)$ for every vertex w of τ. Then for every point $x \in \tau$,

$$\angle(N_\tau, N_{\tilde{x}}\Sigma) \leq \max\left\{\eta(\sqrt{2}\kappa) + \arcsin\left(\kappa \cot\frac{\phi}{2}\right), \eta(2\kappa) + \arcsin\left(\kappa\cot\left(45° - \frac{\phi}{4}\right)\right)\right\}$$

where $\eta(\delta) = \eta_1(\delta)$ in codimension 1, or $\eta(\delta) = \eta_2(\delta)$ in higher codimensions; and ϕ is a "proof parameter" that can be set to any angle in the range $(0°, 60°]$. We recommend choosing $\phi = 49°$ in codimension 1, and $\phi = 48.5°$ in higher codimensions. Figure 2.6 graphs the bound for both cases. We give an alternative version of this bound tailored for restricted Delaunay triangles in an ϵ-sample of Σ. (See Section 2.5.)

Beyond the improved approximation bounds and their extensions to higher dimensions, we think that some of the proof ideas in this chapter are interesting in their own right. Our proof of the

Triangle Normal Lemma is strongly intuitive and reveals a lot about *why* the bound is what it is. Our proofs of the Normal Variation Lemmas exploit properties of medial balls and medial-free balls in ways that allow us to obtain stronger bounds than prior proofs, which were based on integration of the curvature along a path on Σ. These properties also find application in sequel work that improves the sampling bounds needed to guarantee that a triangulation is homeomorphic to an underlying 2-manifold.

In many applications (such as mechanical modeling of stress), the interpolation error in the gradient, $\|\nabla f(p) - \nabla g(p)\|$, is even more important than $|f(p) - g(p)|$. The pointwise gradient interpolation error $\|\nabla f(p) - \nabla g(p)\|$ at the worst point p in a simplex scales linearly with the size of the simplex, and is very sensitive to the shape of the simplex. An early analysis by Bramble and Zlámal [15] for $\mathbb{R}^2$ seemed to implicate triangles with small angles (near 0°), but a famous paper by Babuška and Aziz [8] vindicated small angles and placed the blame on large angles (near 180°). A triangle's circumradius alone suffices to produce a reasonable rough bound on the pointwise gradient interpolation error over the triangle, but a stronger bound can be obtained by taking into account additional information about the triangle's shape [83]. Similarly, in this chapter we show that a triangle's circumradius alone suffices to produce a reasonable rough bound (2.2) on the normal error, but a stronger bound (2.1) can be obtained by taking into account more information about shape.

2.3 Triangle Normal Lemmas

Given a triangle τ whose vertices lie on a k-manifold Σ, we derive bounds on how well τ's normal space locally approximates the spaces normal to Σ in the vicinity of τ. In this section, we derive a bound on $\angle(N_\tau, N_v\Sigma) = \angle(\operatorname{aff}\tau, T_v\Sigma)$ where v is a vertex of τ. (In codimension 1, we can interpret this as the angle between normal vectors, albeit a nonobtuse angle—we do not distinguish between a vector n_v and its negation $-n_v$.) We first consider surfaces embedded in $\mathbb{R}^3$, then we show that the same bound applies to k-manifolds embedded in $\mathbb{R}^d$ for all $d > k \geq 2$. In Section 2.5, we give a bound on $\angle(N_\tau, N_{\tilde{x}}\Sigma) = \angle(\operatorname{aff}\tau, T_{\tilde{x}}\Sigma)$ applicable to every point $x \in \tau$, not just at the vertices. Hence, it applies to the normal spaces of all the points in $\nu(\tau)$. Note that in the lemma, each occurrence of $\operatorname{ebs}(v)$ can be replaced by $\operatorname{lfs}(v)$, as $\operatorname{lfs}(v) \leq \operatorname{ebs}(v)$.

Lemma 1 (Triangle Normal Lemma for $\mathbb{R}^3$)**.** *Let Σ be a smooth 2-manifold without boundary embedded in $\mathbb{R}^3$. Let τ be a triangle whose vertices lie on Σ. Let R be τ's circumradius. Let v be a vertex of τ and let ϕ be τ's plane angle at v. Then*

$$\angle(N_\tau, N_v\Sigma) = \angle(\operatorname{aff}\tau, T_v\Sigma) \leq \arcsin\left(\frac{R}{\operatorname{ebs}(v)} \max\left\{\cot\frac{\phi}{2}, 1\right\}\right).$$

(Note that the argument $\cot\frac{\phi}{2}$ *dominates if ϕ is acute and the argument* 1 *dominates if ϕ is obtuse.) In particular, if v is the vertex at τ's largest plane angle (so $\phi \geq 60°$) and $R < \operatorname{ebs}(v)/\sqrt{3} \doteq 0.577\operatorname{ebs}(v)$, then*

$$\angle(N_\tau, N_v\Sigma) = \angle(\operatorname{aff}\tau, T_v\Sigma) \leq \arcsin\frac{\sqrt{3}R}{\operatorname{ebs}(v)}.$$

Proof. Let $\theta = \angle(N_\tau, N_v\Sigma)$. Consider the two balls of radius $\mathrm{ebs}(v)$ tangent to Σ at v. The plane aff τ intersects these two balls in two circles of radius $\rho = \mathrm{ebs}(v)\sin\theta$, as Figure 2.7 shows. We consider these two circles C_1 and C_2 in the plane aff τ. Notice that since C_1 and C_2 are cross sections of surface-free balls, their insides are surface-free. In particular, u and w cannot lie strictly inside C_1 or C_2. We will use this fact to establish a relationship between the radius ρ of these circles and the circumradius R of τ.

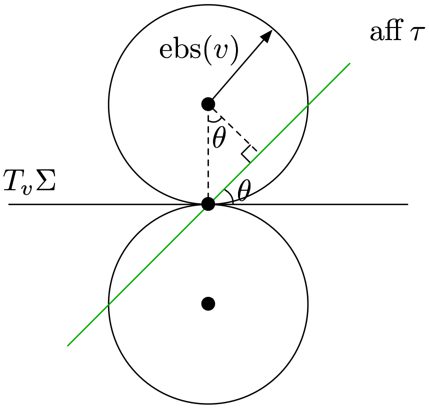

Figure 2.7: The affine hull aff τ intersects the surface-free balls of radius $\mathrm{ebs}(v)$ in two circles of radius $\mathrm{ebs}(v)\sin\theta$.

Let c_1 and c_2 be the centers of C_1 and C_2, respectively. Imagine that as θ increases, and aff τ tilts further, C_1 grows in the direction $v\vec{c}_1$ while remaining in contact with v, and C_2 grows in the opposite direction. We distinguish two cases: (1) either $v\vec{c}_1$ or $v\vec{c}_2$ points into τ or (2) both $v\vec{c}_1$ and $v\vec{c}_2$ point to the exterior of τ. See Figures 2.8 and 2.9.

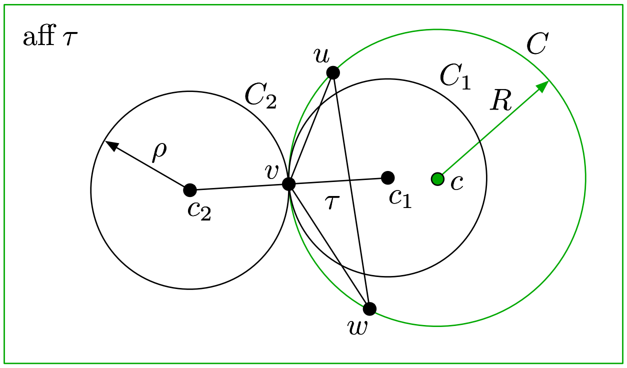

Figure 2.8: Case 1, where one of the two circles grows into the interior of τ. In this case, the radius of C_1 is at most R.

Let $\tau = \triangle uvw$. Let C be the circumcircle of τ in the plane aff τ, and let c be the center of C. In case 1, illustrated in Figure 2.8, one of vc_1 or vc_2 points into τ; suppose it is vc_1. C_1 cannot grow indefinitely; eventually it intersects u or w. The maximum angle is achieved when $C_1 = C$, whereupon u and w prevent further growth. Thus $R \geq \rho = \text{ebs}(v) \sin\theta$ which implies that $\theta \leq \arcsin \frac{R}{\text{ebs}(v)}$.

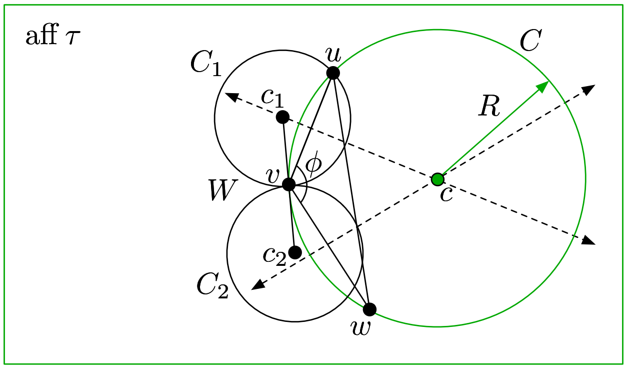

Figure 2.9: Case 2, where both circles grow into the exterior of τ. In this case, the bound depends on the angle ϕ at v.

In case 2, the line segment c_1c_2 does not intersect τ except at v, as Figure 2.9 shows. The bisectors of vu and vw divide the plane into four wedges with apex c; let W be the closed wedge that contains v. As vu and vw meet at v at an angle ϕ, the wedge angle where the bisectors meet at c is $180° - \phi$, as illustrated in Figure 2.10.

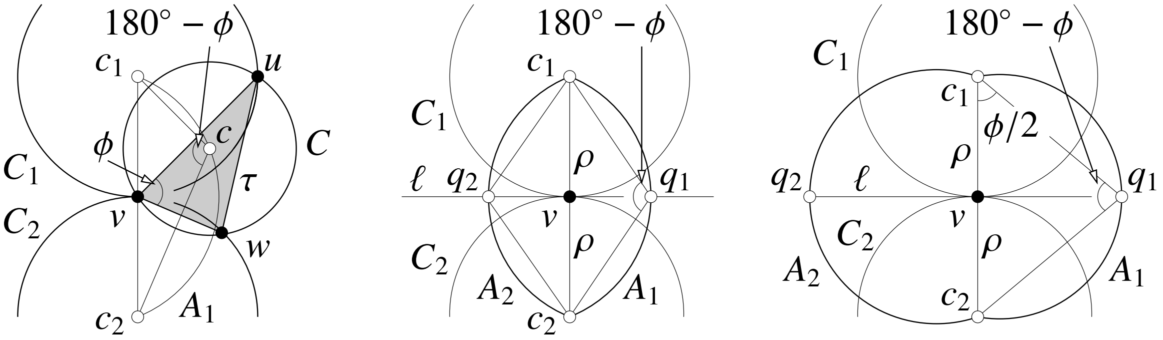

Figure 2.10: Left: the triangle angle of ϕ induces a wedge angle of $180° - \phi$. Center: the circumcenter c cannot lie inside the region enclosed by arcs A_1 and A_2, here illustrated for an acute ϕ. Right: For an obtuse ϕ.

As u is not inside the circle C_1, $|uc_1| \geq |vc_1|$. Similarly, $|wc_1| \geq |vc_1|$. It follows that $c_1 \in W$. Similarly, $c_2 \in W$. Therefore, $\angle c_1cc_2 \leq 180° - \phi$. By circle geometry, this inequality implies that we can draw two circular arcs with endpoints c_1 and c_2 such that c cannot be strictly inside the region enclosed by the arcs. Specifically, let ℓ be the line that bisects c_1c_2. Let q_1 and q_2 be the two distinct points on ℓ such that $\angle c_1q_1c_2 = 180° - \phi$ and $\angle c_1q_2c_2 = 180° - \phi$, as illustrated in Figure 2.10. Both of these angles are bisected by ℓ; that is, $\angle c_iq_jv = 90° - \phi/2$ for $i \in \{1, 2\}$, $j \in \{1, 2\}$. Thus we have four similar right triangles adjoining v of the form $\triangle c_ivq_j$ with $\angle q_jc_iv = \phi/2$.

Observe that $|vc_1| = |vc_2| = \rho = \mathrm{ebs}(v)\sin\theta$, hence $|vq_1| = |vq_2| = \rho\tan(\phi/2)$. Consider the unique circular arc A_1 having endpoints c_1 and c_2 and passing through q_1, and its mirror image arc A_2 passing through q_2, as illustrated. By circle geometry, for every point q on A_1 or A_2 (except c_1 or c_2), $\angle c_1qc_2 = 180° - \phi$, and for every point q enclosed between the two arcs, $\angle c_1qc_2 > 180° - \phi$. It follows that the circumcenter c cannot lie in the region enclosed by A_1 and A_2.

As $\sin\theta \leq \rho/\mathrm{ebs}(v)$, our goal is to determine the maximum possible value of ρ for a fixed value of R. Equivalently, we wish to determine the minimum value of $R = |vc|$ for a fixed ρ. In other words, with ρ fixed, what is the closest that c can get to v? If $\phi \leq 90°$, then the distance $|vc|$ is minimized for $c = q_1$ or $c = q_2$ (see Figure 2.10, center), in which case $R = |vq_1| = \rho\tan(\phi/2)$. If $\phi \geq 90°$, then $|vc|$ is minimized for $c = c_1$ or $c = c_2$ (see Figure 2.10, right), in which case $r = |vc_1| = \rho$. It follows that $r \geq \rho\min\{\tan(\phi/2), 1\}$, hence $\sin\theta \leq \rho/\mathrm{ebs}(v) \leq R\max\{\cot(\phi/2), 1\}/\mathrm{ebs}(v)$. □

Compare Lemma 1 with two prior versions of the Triangle Normal Lemma. The following lemma gives the best known bound, which was proven by Amenta, Choi, Dey, and Leekha [6]. (The derivation can also be found in Dey [28] and Cheng et al. [20].)

Lemma 2. *Let Σ be a smooth 2-manifold without boundary embedded in $\mathbb{R}^3$. Let τ be a triangle whose vertices lie on Σ. Let R be τ's circumradius. Let v be the vertex of τ at τ's largest plane angle. If $R \leq 0.433\,\mathrm{lfs}(v)$, then*

$$\angle(N_\tau, N_v\Sigma) = \angle(\mathrm{aff}\,\tau, T_v\Sigma) \leq \arcsin\left(\frac{R}{\mathrm{lfs}(v)}\right) + \arcsin\left(\frac{2}{\sqrt{3}}\sin\left(2\arcsin\left(\frac{R}{\mathrm{lfs}(v)}\right)\right)\right).$$

The year before, Cheng, Dey, Edelsbrunner, and Sullivan [18] derived a stronger bound of $\arcsin\frac{2R}{\mathrm{lfs}(v)}$, but it seems to have escaped notice. All three bounds are plotted in Figure 2.5 (left). Lemma 1 improves upon both prior results in three ways: it is tighter for the case covered by Lemma 2 (improving the Cheng et al. bound by a factor of 1.15 and the Amenta et al. bound by a factor of 1.91 for small values of $R/\mathrm{lfs}(v)$), it applies to any vertex v of τ, and it takes into account τ's angle at v.

Lemma 1 extends straightforwardly to higher-dimensional manifolds embedded in higher-dimensional Euclidean spaces (but not to higher-dimensional simplices). Given a triangle τ whose vertices lie on a k-manifold $\Sigma \subset \mathbb{R}^d$, we wish to know the worst-case angle deviation $\angle(\mathrm{aff}\,\tau, T_v\Sigma)$ between τ's affine hull and the tangent space at a vertex v of τ.

Lemma 3 (Triangle Normal Lemma for $\mathbb{R}^d$)**.** *Let Σ be a smooth k-manifold without boundary embedded in $\mathbb{R}^d$, with $k \geq 2$. Let τ be a triangle whose vertices lie on Σ. Let R be τ's circumradius. Let v be a vertex of τ and let ϕ be τ's plane angle at v. Then*

$$\angle(N_\tau, N_v\Sigma) = \angle(\mathrm{aff}\,\tau, T_v\Sigma) \leq \arcsin\left(\frac{R}{\mathrm{ebs}(v)}\max\left\{\cot\frac{\phi}{2}, 1\right\}\right).$$

Proof. The dimension of $N_v\Sigma$ is less than or equal to the dimension of N_τ (which is $d - 2$), so by definition,

$$\angle(N_\tau, N_v\Sigma) = \max_{\ell_v \subset N_v\Sigma} \min_{\ell_N \subset N_\tau} \angle(\ell_N, \ell_v)$$

where ℓ_v and ℓ_N are lines. Let $\ell_v \subset N_v\Sigma$ and $\ell_N \subset N_\tau$ be lines such that $\angle(N_\tau, N_v\Sigma) = \angle(\ell_N, \ell_v)$, translated so they pass through v (without loss of generality). If $\angle(\ell_N, \ell_v) = 0$ the result follows immediately, so suppose that $\angle(\ell_N, \ell_v) > 0$. Let Π be the plane (2-flat) that includes both ℓ_v and ℓ_N. Let $\ell_\tau \subset \Pi$ be the line through v perpendicular to ℓ_N in Π. As ℓ_N is chosen from the flat N_τ to minimize its angle with ℓ_v, the line ℓ_τ is orthogonal to N_τ, and therefore ℓ_τ lies in the complementary flat aff τ. Let $\Xi \subset \mathbb{R}^d$ be the unique 3-flat that includes τ and ℓ_N. As Ξ includes aff τ, $\ell_\tau \subset \Xi$; and as Ξ also includes ℓ_N, $\Pi \subset \Xi$, hence $\ell_v \subset \Xi$.

We reiterate the proof of Lemma 1 to bound $\angle(\ell_N, \ell_v)$, with Ξ replacing $\mathbb{R}^3$ and ℓ_v replacing $N_v\Sigma$ in the proof. The proof of Lemma 1 relies entirely on the fact that τ's vertices cannot be inside the two open balls of radius ebs(v) that are centered on ℓ_v and touching v. In the present setting in $\mathbb{R}^d$, every open ball of radius ebs(v) tangent to Σ at v is surface-free; two of those balls have centers on ℓ_v. The intersections of these balls with Ξ are surface-free 3-balls of radius ebs(v), so the constraints harnessed by the proof of Lemma 1 hold in the subspace Ξ. Therefore, the bound of Lemma 1 holds for k-manifolds in $\mathbb{R}^d$ as well. □

2.4 Normal Variation Lemmas

Recall that, given two nearby points $p, q \in \Sigma$, we seek an upper bound on the *normal variation*, the angle $\angle(n_p, n_q)$ separating their normal vectors (in codimension 1) or the angle $\angle(N_p\Sigma, N_q\Sigma)$ separating their normal spaces (in codimension 2 or higher).

Lemma 4 (Normal Variation Lemma for Codimension 1)**.** *Let $\Sigma \subset \mathbb{R}^d$ be a bounded, smooth $(d-1)$-manifold without boundary. Consider two points $p, q \in \Sigma$ and let $\delta = |pq|/\mathrm{lfs}(p)$. Let n_p and n_q be outward-directed vectors normal to Σ at p and q, respectively.*

If $\delta < \sqrt{4\sqrt{5} - 8} \doteq 0.9717$, then $\angle(n_p, n_q) \le \eta_1(\delta)$ where

$$\eta_1(\delta) = \arccos\left(1 - \frac{\delta^2}{2\sqrt{1-\delta^2}}\right) \approx \delta + \frac{7}{24}\delta^3 + \frac{123}{640}\delta^5 + \frac{1{,}083}{7{,}168}\delta^7 + O(\delta^9). \tag{2.3}$$

Moreover, if δ_N is the component of δ parallel to p's normal line $N_p\Sigma$—that is, δ_N is the distance from q to the tangent space $T_p\Sigma$ divided by lfs(p)*—we have the bound (which is stronger when $\delta_N \neq 0$)*

$$\angle(n_p, n_q) \le \arccos\left(1 - \frac{\delta^2 - \delta^4/2 - 2\delta_N^2}{\sqrt{(1-\delta^2)\left((2-\delta^2)^2 - 4\delta_N^2\right)}}\right). \tag{2.4}$$

Recall that the right-hand side of Inequality (2.3) is plotted in green in Figure 2.4. Two isocontour plots of the right-hand side of Inequality (2.4) appear in Figure 2.11. In most circumstances where

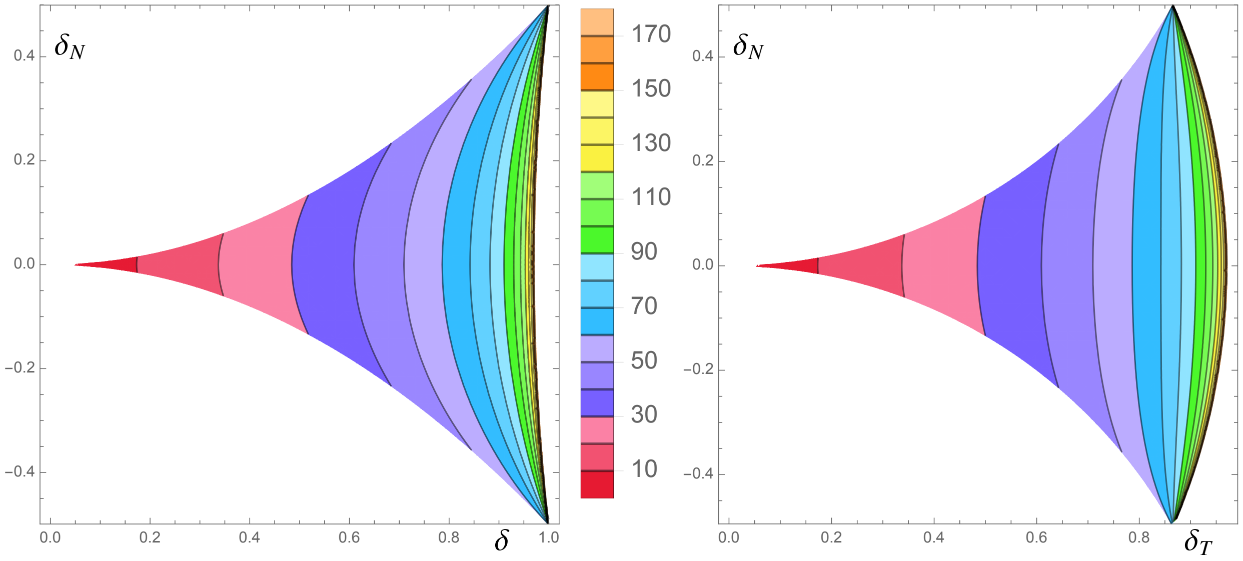

Figure 2.11: Left: The upper bound (in degrees) for $\angle(n_p, n_q)$ as specified by Inequality (2.4), as a function of $\delta = |pq|/\text{lfs}(p)$ (on the horizontal axis) and the normal component δ_N of δ (on the vertical axis); i.e., δ_N is the distance from q to the tangent space $T_p\Sigma$ divided by lfs(p). Right: A similar plot with one change: the horizontal axis is the tangential component δ_T of δ; i.e., the distance from q to the normal line $N_p\Sigma$ divided by lfs(p). This plot reflects the Euclidean geometry of the space, with p at the origin, $T_p\Sigma$ on the horizontal midline, q somewhere in the colored region, and the two surface-free balls of radius lfs(p) (white) blocking q from occupying certain regions (compare with Figure 2.12).

a normal variation lemma is applied, $|pq|$ is known but the normal component δ_N is not. It is clear from the plot on the left that for any given value of δ, the bound (2.4) is weakest at $\delta_N = 0$; this substitution yields the bound (2.3). Hence the green curve in Figure 2.4 also represents the horizontal midline of the isocontour plot.

Proof. Let F be the open ball with center p and radius lfs(p). By the definition of lfs, F does not intersect the medial axis M of Σ. The line $N_p\Sigma$ normal to Σ at p intersects the boundary of F at two opposite poles o and o'. By assumption, $|pq| < \text{lfs}(p)$, so $q \in F$ and the normal line $N_q\Sigma$ intersects the boundary of F at two points z and z'.

Let B and B' be the two open balls of radius lfs(p) tangent to Σ at p, illustrated in Figure 2.12; the centers of these balls are o and o', respectively. Neither ball intersects Σ nor contains q. Let Z be the open ball centered at z with its boundary passing through q, and define Z' likewise with its center at z'. Each of Z and Z' is a subset of a medial ball tangent to Σ at q, so neither ball intersects Σ nor contains p. Without loss of generality, suppose that B' and Z' are enclosed by Σ, whereas B and Z are outside the region enclosed by Σ. Therefore, B is disjoint from Z', and B' is disjoint from Z. (However, B may intersect Z, and B' may intersect Z'.) This property is the key to obtaining a bound on $\angle(n_p, n_q)$.

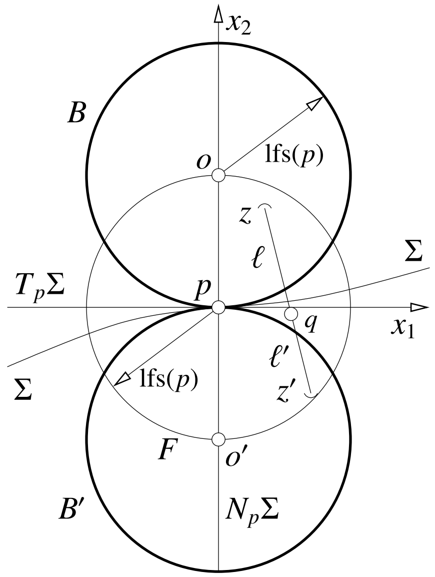

Figure 2.12: The medial-free ball F and surface-free balls B and B' associated with a point $p \in \Sigma$.

We create a d-axis coordinate system with p at the origin. For simplicity, we will scale the coordinate system so that $\mathrm{lfs}(p) = 1$; hence B, B', and F all have radius 1. The x_2-axis is the normal line $N_p\Sigma$, which passes through o, p, and o' and is directed so that $o = (0, 1, 0, \ldots, 0)$, $o' = (0, -1, 0, \ldots, 0)$, and $p = (0, 0, \ldots, 0)$, as illustrated in Figure 2.12. The remaining axes span the tangent space $T_p\Sigma$. We choose an x_1-axis on $T_p\Sigma$ such that its positive branch passes through the orthogonal projection of q onto $T_p\Sigma$; that is, $q_1 \geq 0$ and $q_3 = q_4 = \ldots = q_d = 0$. We choose an x_3-axis on $T_p\Sigma$ such that the normal line $N_q\Sigma$ lies in the x_1-x_2-x_3-space (which is now the affine hull of $N_p\Sigma \cup N_q\Sigma$). Hence, $z_4 = z_5 = \ldots = z_d = 0$ and $z'_4 = z'_5 = \ldots = z'_d = 0$. All the important features of the problem lie on the three-dimensional cross-section of $\mathbb{R}^d$ specified by these three coordinates.

Let $\ell = |qz|$ and $\ell' = |qz'|$ be the radii of Z and Z', respectively. The unit ball F has a diameter e that passes through q (and through the origin p, like all diameters of F). The point q subdivides e into a line segment of length $1 + \|q\|$ and a line segment of length $1 - \|q\|$. As this diameter and the line segment zz' intersect each other at q, they are both chords of a common circle on the boundary of F, illustrated in Figure 2.13. By the well-known Intersecting Chords Theorem,

$$\ell\ell' = (1 + \|q\|)(1 - \|q\|) = 1 - \|q\|^2, \tag{2.5}$$

where $\|q\|^2 = q_1^2 + q_2^2$ (as q's other coordinates are zero). Note that $\|q\|$ is the distance from p to q.

The balls Z and B' (with centers z and o' and radii ℓ and 1) are disjoint and z lies on the unit

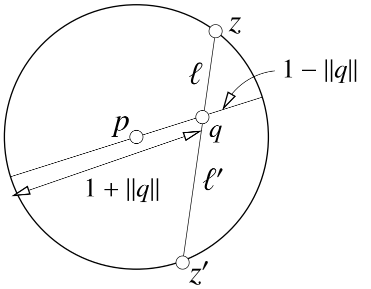

Figure 2.13: The Intersecting Chords Theorem: $\ell\ell' = (1 + \|q\|)\,(1 - \|q\|)$.

sphere, so

$$\begin{aligned} \ell + 1 &\leq |zo'| \\ &= \sqrt{z_1^2 + (z_2 + 1)^2 + z_3^2} \\ &= \sqrt{2 + 2z_2}. \end{aligned} \tag{2.6}$$

Symmetrically, Z' and B are disjoint, so

$$\ell' + 1 \leq \sqrt{2 - 2z_2'}. \tag{2.7}$$

If one of the inequalities (2.6) or (2.7) holds with equality, we call this event a *tangency*. A tangency between Z and B' implies that

$$z_2 = \frac{(\ell + 1)^2}{2} - 1, \tag{2.8}$$

whereas a tangency between Z' and B implies that

$$z_2' = 1 - \frac{(\ell' + 1)^2}{2}. \tag{2.9}$$

Our goal is to find an upper bound on $\angle(n_p, n_q)$. This angle is the tilt of the line segment zq relative to the x_2-axis, so

$$\cos\angle(n_p, n_q) = \frac{z_2 - q_2}{|zq|} = \frac{z_2 - q_2}{\ell}.$$

To find a bound, we seek to determine the configuration(s) in which the angle is maximized—hence, the cosine is minimized—subject to Inequalities (2.6) and (2.7). We will see that the maximum is obtained when both inequalities hold with equality, a configuration we call a *dual tangency*, illustrated in Figure 2.14.

In a configuration where neither tangency is engaged (i.e., both inequalities are strict), we can increase $\angle(n_p, n_q)$ and decrease its cosine by freely tilting the line segment zz' while maintaining

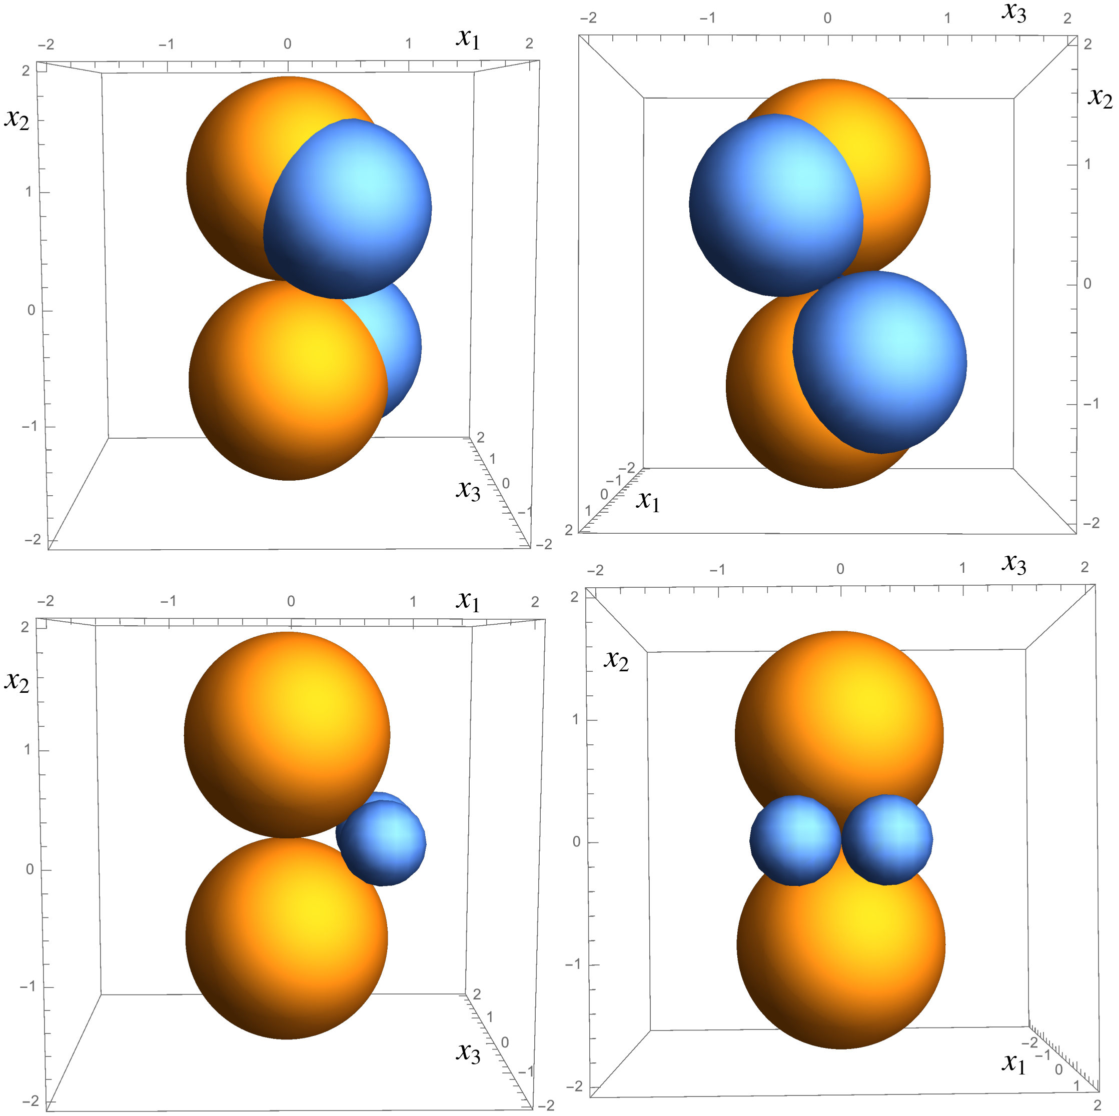

Figure 2.14: Dual tangency configurations for $\delta = 0.5$ (top two images) and $\delta = 0.9101$ (bottom two images). In the former configuration, $\angle(n_p, n_q) \doteq 31.17°$, and in the latter configuration, $\angle(n_p, n_q) \doteq 90°$. The orange balls are B and B', with p at their point of tangency, and the blue balls are Z and Z', with q at their point of tangency. The manifold Σ passes through p and q but does not intersect the interior of any of these balls.

the constraints that zz' passes through q, and both z and z' lie on the boundary of F. (Note that in our coordinate system, q, p, B, B', F, and n_p are all fixed, but we can adjust n_q subject to the inequalities.) Therefore, if the maximum possible angle is not 180°, a configuration that maximizes the angle must engage at least one tangency. As Z and Z' play symmetric roles, we can assume without loss of generality that Z is tangent to B' and Equation (2.8) holds, giving

$$\cos \angle(n_p, n_q) = 1 + \frac{\ell^2 - 1 - 2q_2}{2\ell}. \tag{2.10}$$

The derivative $\frac{\partial}{\partial \ell} \cos \angle(n_p, n_q) = (\ell^2 + 1 + 2q_2)/(2\ell^2)$ is positive for all $q_2 \geq -1/2$; we have $q_2 \in (-1/2, 1/2)$ because $q \in F$, $q \notin B'$, and $q \notin B$. Therefore, the cosine (2.10) increases monotonically with ℓ. We see from Equation (2.5) that ℓ increases monotonically as ℓ' decreases. Inequality (2.7) places an upper bound on ℓ', which together with (2.5) places a lower bound on ℓ, which places a lower bound on the cosine (2.10) and an upper bound on the angle $\angle(n_p, n_q)$ itself. A configuration attains this upper bound on $\angle(n_p, n_q)$ when Inequality (2.7) holds with equality—in a dual tangency, where Z' is tangent to B in addition to Z being tangent to B',

A dual tangency uniquely determines the values of ℓ and ℓ'. As $q \in zz'$, we can write

$$\ell(z'_2 - q_2) = \ell'(q_2 - z_2). \tag{2.11}$$

The identities (2.5), (2.8), (2.9), and (2.11) form a system of four (nonlinear) equations in the four variables ℓ, ℓ', z_2, and z'_2. According to Mathematica (and verified by substitution), these equations are simultaneously satisfied by

$$\ell = \sqrt{\frac{(1 - \|q\|^2)\,(2 + 2q_2 - \|q\|^2)}{2 - 2q_2 - \|q\|^2}} \quad \text{and} \quad \ell' = \sqrt{\frac{(1 - \|q\|^2)\,(2 - 2q_2 - \|q\|^2)}{2 + 2q_2 - \|q\|^2}}. \tag{2.12}$$

As this configuration places a lower bound on ℓ, substituting the identity (2.12) into (2.10) shows that

$$\cos \angle(n_p, n_q) \geq 1 - \frac{\|q\|^2 - \|q\|^4/2 - 2q_2^2}{\sqrt{(1 - \|q\|^2)\left((2 - \|q\|^2)^2 - 4q_2^2\right)}}. \tag{2.13}$$

Recall the parameter $\delta = |pq|/\mathrm{lfs}(p)$. As we chose and scaled our coordinate system so that p is the origin and $\mathrm{lfs}(p) = 1$, $\|q\| = \delta$ and $q_2 = \delta_N$. Inequality (2.4) follows.

This expression provides a strong upper bound when the value of q_2 (the distance from q to $T_p\Sigma$) is known, but q_2 is not usually available in circumstances where the Normal Variation Lemma is invoked. To find a bound independent of q_2, we seek the value of $q_2 \in (-\|q\|^2/2, \|q\|^2/2)$ that minimizes the right-hand side of (2.13). The left plot in Figure 2.11 makes it clear that for all $\|q\| < 1$, this value is $q_2 = 0$. To verify this formally, observe that (2.13) is symmetric about $q_2 = 0$ (as it is a function of q_2^2) and

$$\frac{\partial}{\partial q_2} \cos \angle(n_p, n_q) = 2q_2 \frac{3(1 - \|q\|^2)^2 + 4(1 - \|q\|^2) + (1 - 4q_2^2)}{\sqrt{1 - \|q\|^2}\left((2 - \|q\|^2)^2 - 4q_2^2\right)^{3/2}}.$$

The numerator and denominator are positive for all $\|q\| < 1$ and $q_2 \in (-0.5, 0.5)$, so the derivative is zero at $q_2 = 0$, positive for $q_2 > 0$, and negative for $q_2 < 0$, showing that the cosine is minimized at $q_2 = 0$. Setting $q_2 = 0$ shows that

$$\cos\angle(n_p, n_q) \geq 1 - \frac{\|q\|^2}{2\sqrt{1 - \|q\|^2}},$$

proving Inequality (2.3). □

We conjecture (but are not certain) that Inequality (2.3) is sharp: for every legal δ, there exists a surface Σ and points $p, q \in \Sigma$ for which the bound holds with equality. Proving this conjecture would entail finding a surface Σ that is compatible with the four balls B, B', Z, and Z' in the dual tangency described in the proof of Lemma 4 and illustrated in Figure 2.14—meaning that Σ intersects none of the four balls but passes through the four points of tangency p, q, z, and z'—such that no point of Σ's medial axis lies in the ball F.

Figure 2.14 reveals that in the worst-case configuration, n_q is tilted along the x_3-axis (so $z_3 = -z'_3 \neq 0$), but not along the x_1-axis (i.e., $z_1 = z'_1 = q_1$). In other words, Σ undergoes a helical twisting as one walks from p to q. By contrast, a tilt along the x_1-axis cannot be as large.

The proof of the Normal Variation Lemma for higher codimensions is similar in many respects, but it takes a different turn because adding an extra dimension to the normal space enables a novel configuration (not possible in codimension 1) such that the largest angle no longer occurs when $q \in T_p\Sigma$.

Lemma 5 (Normal Variation Lemma for Codimension 2 and Higher)**.** *Let $\Sigma \subset \mathbb{R}^d$ be a bounded, smooth k-manifold without boundary for any $k < d$. Consider two points $p, q \in \Sigma$ and let $\delta = |pq|/\mathrm{lfs}(p)$.*

If $\delta < \sqrt{\left(\sqrt{5} - 1\right)/2} \doteq 0.7861$, then $\angle(N_p\Sigma, N_q\Sigma) = \angle(T_p\Sigma, T_q\Sigma) \leq \eta_2(\delta)$ where

$$\eta_2(\delta) = \arccos\sqrt{1 - \frac{\delta^2}{\sqrt{1 - \delta^2}}} \approx \delta + \frac{5}{12}\delta^3 + \frac{57}{160}\delta^5 + \frac{327}{896}\delta^7 + O(\delta^9). \tag{2.14}$$

Moreover, if δ_N is the component of δ parallel to p's normal space $N_p\Sigma$—that is, δ_N is the distance from q to the tangent space $T_p\Sigma$ divided by lfs(p)*—we have the (stronger) bound*

$$\angle(N_p\Sigma, N_q\Sigma) \leq \arccos\sqrt{\left(1 - \frac{\delta^2}{2\sqrt{1 - \delta^2}}\right)^2 - \frac{\delta_N^2}{1 - \delta^2}}. \tag{2.15}$$

*In the special case where $q \in T_p\Sigma$ (that is, $\delta_N = 0$), this bound reduces to the codimension-*1 *bound $\eta_1(\delta)$ from Lemma 4.*

An isocontour plot of the right-hand side of Inequality (2.15) appears in Figure 2.15. For any given value of δ, the bound (2.15) is weakest along the upper (or lower) boundary of the plot, at $\delta_N = \delta^2/2$; this substitution yields the bound (2.14). The upper boundary is also plotted as the red

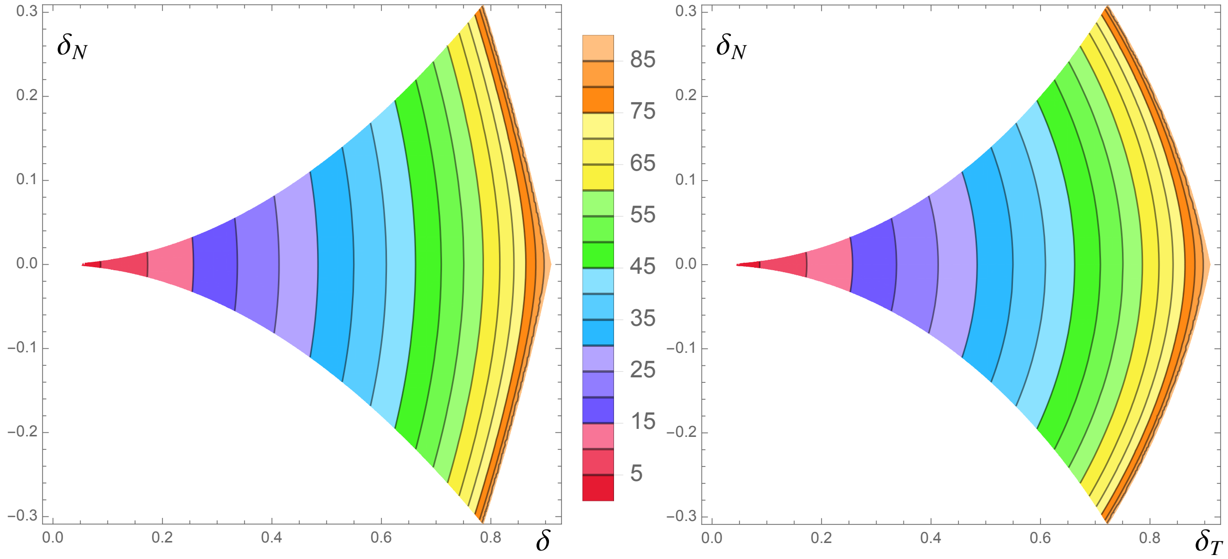

Figure 2.15: Left: Upper bound in degrees for $\angle(N_p\Sigma, N_q\Sigma)$ as specified by Inequality (2.15), as a function of $\delta = |pq|/\mathrm{lfs}(p)$ (on the horizontal axis) and the normal component δ_N of δ (on the vertical axis); i.e., δ_N is the signed distance from q to the tangent space $T_p\Sigma$ divided by $\mathrm{lfs}(p)$. Right: A similar plot with one change: the horizontal axis is the tangential component δ_T of δ; i.e., the distance from q to the normal space $N_p\Sigma$ divided by $\mathrm{lfs}(p)$. Hence, this plot reflects the Euclidean geometry of the space, with p at the origin, $T_p\Sigma$ on the horizontal midline, q somewhere in the colored region, and the smallest possible medial torus (white) blocking q from certain regions.

curve in Figure 2.4. Interestingly, the horizontal midline of this plot is the green curve in Figure 2.4: when $\delta_N = 0$, the symmetry of the configuration yields the codimension-1 bound $\eta_1(\delta)$. The bound gets worse from there as δ_N increases.

We are certain that this bound can be tightened for larger values of δ_N (but not for $\delta_N = 0$), but we have not been able to derive a better explicit bound. It would be nice if the codimension 1 bound held for all δ_N, but we think it very unlikely; we know a configuration in $\mathbb{R}^4$ that defies the codimension 1 bound and which we think (but don't know for sure) can be realized by a 2-manifold fitting the specified constraints.

Proof. Let F be the open ball with center p and radius $\mathrm{lfs}(p)$. F does not intersect the medial axis. As in the proof of Lemma 4, we choose a coordinate system with p at the origin and scale the coordinate system so that $\mathrm{lfs}(p) = 1$, so F is the unit ball centered at the origin.

Let $\dot{\mathbb{B}}$ be the intersection of p's normal space $N_p\Sigma$ with the unit hypersphere ∂F (the boundary of F); $\dot{\mathbb{B}}$ is a unit $(d-k-1)$-sphere. For every point $c \in \dot{\mathbb{B}}$, the open unit ball with center c is tangent to Σ at p and does not intersect Σ. Let $\mathbb{B}$ be the union of these (infinitely many) open unit balls (which constitute all the unit balls tangent to Σ at p). The boundary of $\mathbb{B}$ is a torus with inner radius zero (a horn torus). We call $\mathbb{B}$ itself the (open) *solid torus* and $\dot{\mathbb{B}}$ the *torus skeleton*. Geometrically, $\mathbb{B}$ is the Minkowski sum of $\dot{\mathbb{B}}$ and an open d-ball. Topologically, $\mathbb{B}$ is the d-dimensional product of a $(d-k-1)$-sphere and an open $(k+1)$-ball. Like the balls it is composed of, $\mathbb{B}$ does not intersect Σ nor contain q.

By assumption, $|pq| < \mathrm{lfs}(p)$, so $q \in F$ and q's normal space $N_q\Sigma$ intersects ∂F in a $(d-k-1)$-sphere S (like $\dot{\mathbb{B}}$, but smaller). Consider an open ball Z with center $z \in S$ such that Z's boundary passes through q. Z is a subset of a medial ball tangent to Σ at q, so Z does not intersect Σ nor contain p.

The key property for obtaining a bound is that Z cannot intersect every open unit ball centered on $\dot{\mathbb{B}}$. If it did, then it would effectively block the hole in the solid torus $\mathbb{B}$, so that Σ cannot thread through $\mathbb{B}$ at q without somewhere intersecting Z or $\mathbb{B}$. This property applies to every ball Z centered on S and just touching q. To obtain a tractable proof, we focus on two particular balls that help determine the angle $\angle(N_p\Sigma, N_q\Sigma)$. (Unfortunately, these two balls do not suffice to give a sharp bound, but we have not been able to derive better closed-form bounds that take advantage of the other balls.)

We choose a d-axis coordinate system with p at the origin such that the x_1-axis lies on p's tangent space $T_p\Sigma$, the x_2-axis lies on p's normal space $N_p\Sigma$, and q lies in the upper right quadrant of the x_1-x_2-plane; that is, $q_1 > 0$, $q_2 \geq 0$, and $q_3 = q_4 = \ldots = q_d = 0$. Each remaining axis lies in $T_p\Sigma$ or $N_p\Sigma$, so every axis can be categorized as tangential or normal with respect to p. Let z_T^2 be the sum of squares of the tangential components of z except z_1, and let z_N^2 be the sum of squares of the normal components of z except z_2; thus $\|z\|^2 = z_1^2 + z_2^2 + z_T^2 + z_N^2$. (The signs of z_T and z_N are irrelevant.)

By definition, $\angle(N_p\Sigma, N_q\Sigma) = \max_{\ell_q \subset N_q\Sigma} \min_{\ell_p \subset N_p\Sigma} \angle(\ell_p, \ell_q)$. Let $\ell_q \subset N_q\Sigma$ be a line through q that satisfies $\angle(N_p\Sigma, N_q\Sigma) = \angle(N_p\Sigma, \ell_q)$. Let z and z' be the two points where ℓ_q intersects ∂F, and observe that $z, z' \in S$ (as $S = N_q\Sigma \cap \partial F$). Let Z and Z' be the open balls centered on z and z', respectively, with the boundaries of both balls passing through q. Let $\ell = |qz|$ and $\ell' = |qz'|$ be their radii.

As $q_N = 0$, we can determine the angle $\angle(N_p\Sigma, N_q\Sigma)$ from the identity

$$\cos\angle(N_p\Sigma, N_q\Sigma) = \cos\angle(N_p\Sigma, \ell_q) = \frac{\sqrt{(z_2 - q_2)^2 + z_N^2}}{\ell}, \tag{2.16}$$

because the denominator is the length of the line segment qz and the numerator is the length of the projection of qz onto $N_p\Sigma$. To find an upper bound on $\angle(N_p\Sigma, N_q\Sigma)$, we seek a lower bound on the cosine (2.16); to find that, we will search for legal values of z_2, z_N, and ℓ that minimize the right-hand side (i.e., a worst-case configuration). First, we must understand the constraints on these values.

Let o be the point on the torus skeleton $\dot{\mathbb{B}}$ farthest from z. What is the distance $|zo|$? First consider the projection $\bar{z}$ of z onto $N_p\Sigma$. The origin lies between $\bar{z}$ and the farthest point on $\dot{\mathbb{B}}$, so the distance from $\bar{z}$ to the farthest point is $\|\bar{z}\| + 1$. With Pythagoras' Theorem we add the tangential component:

$$\begin{aligned}
|zo|^2 &= |z\bar{z}|^2 + (\|\bar{z}\| + 1)^2 \\
&= z_1^2 + z_T^2 + \left(\sqrt{z_2^2 + z_N^2} + 1\right)^2 \\
&= z_1^2 + z_T^2 + z_2^2 + z_N^2 + 2\sqrt{z_2^2 + z_N^2} + 1 \\
&= 2 + 2\sqrt{z_2^2 + z_N^2}.
\end{aligned}$$

The last step follows because z lies on ∂F.

As Z has radius ℓ and is disjoint from the unit ball centered at o, $\ell + 1 \le |zo|$. We rewrite this constraint as

$$z_2^2 + z_N^2 \ge \left(\frac{(\ell+1)^2}{2} - 1\right)^2. \tag{2.17}$$

If Inequality (2.17) holds with equality, we call this event a *tangency* between Z and $\mathbb{B}$. Likewise, the ball Z' entails the following inequality, and a tangency between Z' and $\mathbb{B}$ means that it holds with equality.

$${z'_2}^2 + {z'_N}^2 \ge \left(\frac{(\ell'+1)^2}{2} - 1\right)^2. \tag{2.18}$$

Recall from the proof of Lemma 4 that, by the Intersecting Chords Theorem, $\ell\ell' = 1 - \|q\|^2$ where $\|q\| = q_1^2 + q_2^2$ is the distance from p to q. As $q \in zz'$, we write two more useful identities:

$$\begin{aligned}
\ell z'_N &= -\ell' z_N, && (2.19)\\
\ell(z'_2 - q_2) &= \ell'(q_2 - z_2). && (2.20)
\end{aligned}$$

Thus we have a system of three equations and two inequalities in six variables: ℓ, ℓ', z_2, z'_2, z_N, and z'_N. Among the multiple solutions of this system, we seek one that minimizes the objective (2.16).

In a configuration where neither tangency is engaged, we can increase $\angle(N_p\Sigma, N_q\Sigma)$ and decrease its cosine (2.16) by freely tilting the line segment zz' while maintaining the constraints that zz' passes through q and $z, z' \in \partial F$. Therefore, if there is a meaningful bound at all, an optimal (i.e, worst-case) configuration must engage at least one tangency. As Z and Z' play symmetric roles, we can assume without loss of generality that Z is tangent to $\mathbb{B}$ and Inequality (2.17) holds with equality. Substituting that identity into (2.16) yields

$$\cos\angle(N_p\Sigma, N_q\Sigma) = \frac{\sqrt{\left(\frac{(\ell+1)^2}{2} - 1\right)^2 + q_2^2 - 2q_2z_2}}{\ell} = \sqrt{\left(1 - \frac{\ell^2 - 1}{2\ell}\right)^2 + \frac{q_2^2 - 2q_2z_2}{\ell^2}}. \tag{2.21}$$

As in the proof of Lemma 5, symmetry will play a role: the “optimal” (i.e., worst-case) solution will turn out to have $\ell = \ell'$. To expose this symmetry, we define a parameter

$$\gamma = \frac{\ell'}{\ell}.$$

By Identities (2.19) and (2.20), we can eliminate the primed variables with the substitutions $\ell' = \gamma\ell$, $z'_N = -\gamma z_N$ and $z'_2 = q_2 + \gamma(q_2 - z_2)$. (A solution with $\gamma = 1$ would imply that $\ell = \ell'$ and $z'_N = -z_N$.) Inequality (2.18) becomes

$$(q_2 + \gamma(q_2 - z_2))^2 + \gamma^2 z_N^2 \geq \left(\frac{(\gamma\ell + 1)^2}{2} - 1\right)^2. \tag{2.22}$$

To eliminate the variable z_N, we multiply Inequality (2.17) by γ^2 (recalling that the inequality is now assumed to be an equality) and subtract Inequality (2.22) (which is still an inequality), giving

$$(2\gamma^2 + 2)q_2 z_2 - (\gamma + 1)^2 q_2^2 \leq \omega \quad \text{where} \quad \omega = \gamma^2\left(\frac{(\ell + 1)^2}{2} - 1\right)^2 - \left(\frac{(\gamma\ell + 1)^2}{2} - 1\right)^2. \tag{2.23}$$

Rearranging, we have

$$q_2 z_2 \leq \frac{(\gamma + 1)^2 q_2^2 + \omega}{2\gamma^2 + 2}. \tag{2.24}$$

Substituting this into (2.21) gives

$$\cos \angle(N_p\Sigma, N_q\Sigma) \geq \sqrt{\left(1 - \frac{\ell^2 - 1}{2\ell}\right)^2 - \frac{2\gamma q_2^2 + \omega}{(\gamma^2 + 1)\,\ell^2}}. \tag{2.25}$$

The right-hand side is a function of γ, ℓ, and the point q. However, the definition $\gamma = \ell'/\ell$ and Equation (2.5) together imply that $\ell = \sqrt{(1 - \|q\|)^2/\gamma}$, so we can write the right-hand side as a function $f(\gamma, q)$. We claim that for all valid q, $f(\gamma, q)$ is minimized at $\gamma = 1$. It is straightforward but tedious (and best done with Mathematica) to verify that $f(\gamma, q) = f(1/\gamma, q)$ and that $\frac{\partial}{\partial\gamma} f(\gamma, q)$ is zero at $\gamma = 1$, positive for $\gamma > 1$, and negative for $\gamma \in (0, 1)$. Specifically, with the abbreviation $\mathring{q} = 1 - \|q\|^2$, we have

$$\frac{\partial}{\partial\gamma} f(\gamma, q) = (\gamma - 1)\frac{(\gamma + 1)^3\mathring{q}^2 + \left((2\gamma^2 + 8\gamma + 2)\mathring{q} + 4\gamma\right)\sqrt{\gamma\mathring{q}}}{4\gamma(\gamma + 1)^2\sqrt{\gamma\mathring{q}}\sqrt{\gamma(1 - 4q_2^2) + 2\gamma\mathring{q} + (1 - \gamma + \gamma^2)\mathring{q}^2 + \frac{4\left((\gamma^2+1)\mathring{q} - 2\gamma\right)\sqrt{\gamma\mathring{q}}}{\gamma+1}}}.$$

The numerator and denominator are positive for $\gamma > 0$, $\mathring{q} > 0$, and $q_2 \in [0, 0.5]$, so the sign of $\frac{\partial}{\partial\gamma} f$ depends solely on the sign of $\gamma - 1$, confirming that the right-hand side of (2.25) is minimized at $\gamma = 1$.

For $\gamma = 1$, we have $\ell = \ell' = \sqrt{1 - \|q\|^2}$ and $\omega = 0$, so Inequality (2.25) becomes

$$\cos \angle(N_p\Sigma, N_q\Sigma) \geq \sqrt{\left(1 - \frac{\|q\|^2}{2\sqrt{1 - \|q\|^2}}\right)^2 - \frac{q_2^2}{1 - \|q\|^2}}, \tag{2.26}$$

Recall the parameter $\delta = |pq|/\mathrm{lfs}(p)$. As we chose and scaled our coordinate system so that p is the origin and $\mathrm{lfs}(p) = 1$, $\|q\| = \delta$. Inequality (2.15) follows.

Clearly, larger values of q_2^2 make the right-hand side smaller (and the bound weaker). It is smallest when q_2 reaches its maximum allowable value of $\|q\|^2/2$. (This maximum is imposed by the fact that $q \notin \mathbb{B}$.) Hence, the following bound holds for all valid values of q_2.

$$\cos \angle(N_p\Sigma, N_q\Sigma) \geq \sqrt{1 - \frac{\|q\|^2}{\sqrt{1 - \|q\|^2}}},$$

proving Inequality (2.14). □

2.5 Extended Triangle Normal Lemmas

The Triangle Normal Lemmas in Section 2.3 bound $\angle(N_\tau, N_v\Sigma) = \angle(\operatorname{aff} \tau, T_v\Sigma)$ only at a vertex v of τ. Moreover, for vertices where τ has a small plane angle, the bound is poor. Here, we derive a bound on $\angle(N_\tau, N_{\tilde{x}}\Sigma)$ for every $x \in \tau$. The method to accomplish this is not new: a triangle normal lemma establishes a strong bound at a vertex where a triangle has a large plane angle, and a normal variation lemma extends the bound from that anchor over the rest of the triangle. We improve on this formulation a bit by taking advantage of the fact that our Triangle Normal Lemma's bound varies with the plane angle at a vertex: we choose τ's vertex nearest $\tilde{x}$ as the anchor if its angle is at least 49°; otherwise, we choose the vertex with the largest plane angle as the anchor.

We begin with several technical lemmas that help us obtain better bounds. Both lemmas help to constrain where $\tilde{x}$ can lie.

Lemma 6. *Let $\Sigma \subset \mathbb{R}^d$ be a smooth k-manifold. Let τ be a simplex (of any dimension) whose vertices lie on Σ. Let B_τ be a closed d-ball such that $B_\tau \supseteq \tau$ (e.g., τ's smallest enclosing ball or a circumscribing ball). Let r be the radius of B_τ, let v be a vertex of τ, and suppose that $r \leq \operatorname{lfs}(v)/2$. Then for every point $x \in \tau$ that is not a vertex of τ, $\tilde{x} = \nu(x)$ is in the interior of B_τ.*

Proof. Consider a point $x \in \tau$ that is not a vertex of τ. As τ's vertices lie in B_τ, x is in the interior of B_τ. If $\tilde{x} = x$ the lemma follows immediately, so suppose that $\tilde{x} \neq x$ and thus $x \notin \Sigma$. Let B be the open medial ball tangent to Σ at $\tilde{x}$ such that x lies on the line segment $\tilde{x}m$, where m is the center of B, as illustrated in Figure A.2. As B is a medial ball, m lies on the medial axis of Σ.

Recall that B is open and B_τ is closed. If the entire closure of B is in the interior of B_τ, then $\tilde{x}$ is in the interior of B_τ and the lemma follows immediately; so assume it is not. Let C be the intersection of the boundaries of B and B_τ. C cannot be the boundary of B, because we have just assumed that B_τ does not include the closure of B. We show that $C \neq \emptyset$ by ruling out the alternatives: we cannot have B and B_τ disjoint because $x \in B$ and x is in the interior of B_τ; we cannot have $B_\tau \subset B$, as τ's vertices are not in B; and we have already ruled out closure$(B) \subset B_\tau$. Hence C is either a $(d-2)$-sphere (e.g., a circle in $\mathbb{R}^3$) or a single point (with B and B_τ tangent to each other at that point, one inside the other).

If C is a $(d-2)$-sphere, let Π be the unique hyperplane that includes that $(d-2)$-sphere, as illustrated; if C contains a single point, let Π be the hyperplane tangent to B_τ and B at that point. Let $\bar{\Pi}_\tau$ be the closed halfspace bounded by Π that includes $B_\tau \setminus B$, and let Π_τ be the open version of

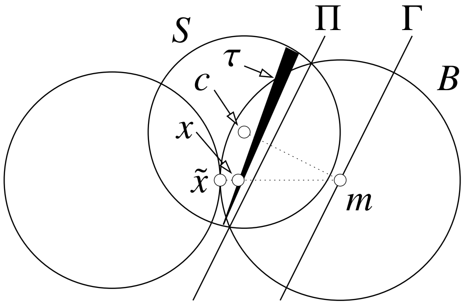

Figure 2.16: For every point $x \in \tau$ except τ's vertices, $\tilde{x}$ is in the interior of B_τ.

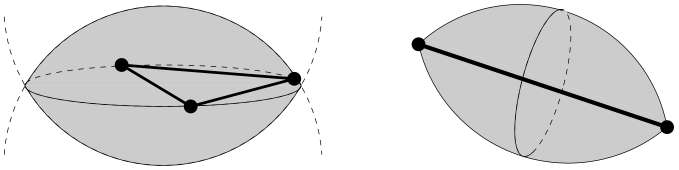

Figure 2.17: For the triangle τ at left, the dark lens-shaped region is the intersection of τ's two enclosing balls of radius $\mathrm{lfs}(v)/2$, where v is any vertex of τ. For every point $x \in \tau$, $\tilde{x}$ lies in this lens. Likewise, for the segment at right, the lemon-shaped region is the intersection of its infinitely many enclosing balls of radius $\mathrm{lfs}(v)/2$; this lemon contains $\tilde{x}$ for every point x on the segment.

the same halfspace. The portion of B in $\bar{\Pi}_\tau$ is in the interior of B_τ, and the portion of B's boundary in Π_τ is in the interior of B_τ. The portion of B_τ in the open halfspace complementary to $\bar{\Pi}_\tau$ is a subset of B. Every vertex of τ lies in B_τ but not in B, hence τ's vertices lie in $\bar{\Pi}_\tau$. Therefore, $\tau \subset \bar{\Pi}_\tau$ and $x \in \bar{\Pi}_\tau$.

By assumption, the radius of B_τ satisfies $r \leq \mathrm{lfs}(v)/2$, so $|vm| \geq \mathrm{lfs}(v) \geq 2r$. As v lies in B_τ and $|vm|$ is at least twice the radius of B_τ, it follows that m is not in the interior of B_τ. But $m \in B$, so $m \notin \bar{\Pi}_\tau$.

Given the facts that x lies on the line segment $m\tilde{x}$, $m \notin \bar{\Pi}_\tau$, $x \in \bar{\Pi}_\tau$, and $\tilde{x} \neq x$, it follows that $\tilde{x} \in \Pi_\tau$. As $\tilde{x}$ is also on B's boundary, $\tilde{x}$ is in the interior of B_τ. □

Lemma 52 implies that $\tilde{x}$ is in *every* ball $B_\tau \supseteq \tau$ with radius $\mathrm{lfs}(v)/2$ (or less). The intersection of these balls, illustrated in Figure 2.17, is typically a narrow region, especially if τ is small. The next lemma also places a restriction on the position of $\tilde{x}$.

Lemma 7. *Let $\Sigma \subset \mathbb{R}^d$ be a smooth k-manifold. Let τ be a simplex (of any dimension) whose vertices lie on Σ. Let r be the min-containment radius of τ (i.e., the radius of τ's smallest enclosing ball). Then for every point $x \in \tau$, the distance from $\tilde{x}$ to the nearest vertex of τ is at most $\sqrt{2}r$. Moreover, if $r < \mathrm{ebs}(\tilde{x})$, the distance from $\tilde{x}$ to the nearest vertex of τ is at most*

$$\sqrt{2\,\mathrm{ebs}(\tilde{x})\left(\mathrm{ebs}(\tilde{x}) - \sqrt{\mathrm{ebs}(\tilde{x})^2 - r^2}\right)} \in [r,\ \sqrt{2}r). \qquad (2.27)$$

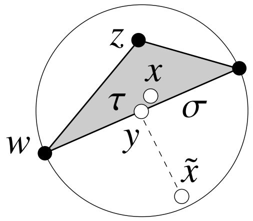

Figure 2.18: Given a simplex τ with min-containment radius r and a point $x \in \tau$, the distance from $\tilde{x}$ to the nearest vertex of τ is at most $\sqrt{2}r$.

Proof. Let $y \in \tau$ be the point nearest $\tilde{x}$ on τ. As x is also on τ, $|y\tilde{x}| \leq |x\tilde{x}|$. Let σ be the unique face of τ (i.e., a vertex, edge, triangle, etc.) whose relative interior contains y. Observe that the line segment $y\tilde{x}$ is orthogonal to σ, as Figure 2.18 illustrates. (If σ is a vertex, it is a trivial "orthogonality.") Let w be the vertex of σ nearest y; $y\tilde{x}$ is orthogonal to yw. By Pythagoras' Theorem, $|w\tilde{x}|^2 = |yw|^2 + |y\tilde{x}|^2 \leq |yw|^2 + |x\tilde{x}|^2$.

As τ's smallest enclosing ball has radius r, $|yw| \leq r$. Likewise, let z be the vertex of τ nearest x; then $|xz| \leq r$. As z lies on Σ and $\tilde{x}$ is the point nearest x on Σ, $|x\tilde{x}| \leq |xz| \leq r$. Hence $|w\tilde{x}| \leq \sqrt{r^2 + r^2} = \sqrt{2}r$, and the distance from $\tilde{x}$ to the nearest vertex of τ (which may or may not be w) is at most $\sqrt{2}r$ as claimed.

Alternatively, if $r < \text{ebs}(\tilde{x})$, we can substitute the bound for $|x\tilde{x}|$ from the Surface Interpolation Lemma [57], yielding the bound $\sqrt{r^2 + \left(\text{ebs}(\tilde{x}) - \sqrt{\text{ebs}(\tilde{x})^2 - r^2}\right)^2}$, which is equal to (2.27). □

This brings us to the main result of this section.

Lemma 8 (Extended Triangle Normal Lemma). *Let Σ be a bounded k-manifold without boundary in $\mathbb{R}^d$ with $k \geq 2$. Let $\tau = \triangle vv'v''$ be a triangle whose vertices lie on Σ. Let R be τ's circumradius. Suppose that $R \leq \kappa\,\text{lfs}(v)$, $R \leq \kappa\,\text{lfs}(v')$, and $R \leq \kappa\,\text{lfs}(v'')$ for some $\kappa \leq 1/2$. Let x be any point on τ, and let $\tilde{x}$ be the point nearest x on Σ. Then for any angle $\phi \in (0°, 60°]$,*

$$\angle(N_\tau, N_{\tilde{x}}\Sigma) \leq \max\left\{\eta(\sqrt{2}\kappa) + \arcsin\left(\kappa\cot\frac{\phi}{2}\right), \eta(2\kappa) + \arcsin\left(\kappa\cot\left(45° - \frac{\phi}{4}\right)\right)\right\}, \qquad (2.28)$$

where $\eta(\delta) = \eta_1(\delta)$ as defined in Lemma 1 if $d - k = 1$, or $\eta(\delta) = \eta_2(\delta)$ as defined in Lemma 3 if $d - k \geq 2$.

Lemma 8 is unusual because it has a parameter ϕ; the right-hand side of Inequality (2.28) varies a bit with ϕ. The parameter ϕ is a threshold that determines which vertex of τ is used as an anchor. In codimension 1, a good choice of ϕ is 49°, because it balances the two expressions in (2.28) reasonably well and delivers a bound below 90° over the range $\kappa \in [0, 0.3734]$. For a specific value of κ, one can tune ϕ to obtain a slightly better bound, but the improvement is marginal. In codimension 2 or greater, the bound (2.28) is weaker because η_2 is weaker than η_1. A good choice is $\phi = 48.5°$, which delivers a bound below 90° over the range $\kappa \in [0, 0.3527]$. Figure 2.6 graphs the bound (2.28) both for codimension 1 and for higher codimensions.

Proof. Suppose without loss of generality that v is the vertex of τ nearest $\tilde{x}$. Let $w \in \{v, v', v''\}$ be the vertex at τ's largest plane angle. Let B_τ be τ's smallest enclosing ball and observe that its radius is $r \leq R \leq \operatorname{lfs}(v)/2$. By Lemma 52, $\tilde{x} \in B_\tau$, so $|w\tilde{x}| \leq 2r \leq 2\kappa \operatorname{lfs}(w)$. By Lemma 7, $|v\tilde{x}| \leq \sqrt{2}r \leq \sqrt{2}\kappa \operatorname{lfs}(v)$. By the Normal Variation Lemma, $\angle(N_w\Sigma, N_{\tilde{x}}\Sigma) \leq \eta(2\kappa)$ and $\angle(N_v\Sigma, N_{\tilde{x}}\Sigma) \leq \eta(\sqrt{2}\kappa)$.

If τ's plane angle at the vertex v is ϕ or greater, then by the Triangle Normal Lemma (Lemma 1 or 3), $\sin \angle(N_\tau, N_v\Sigma) \leq \frac{R}{\operatorname{lfs}(v)} \cot \frac{\phi}{2} \leq \kappa \cot \frac{\phi}{2}$. Then $\angle(N_\tau, N_{\tilde{x}}\Sigma) \leq \angle(N_{\tilde{x}}\Sigma, N_v\Sigma) + \angle(N_\tau, N_v\Sigma) \leq \eta(\sqrt{2}\kappa) + \arcsin(\kappa \cot \frac{\phi}{2})$.

Otherwise, τ's plane angle at v is less than ϕ, so τ's plane angle at w (τ's largest plane angle) is greater than $(180° - \phi)/2$. By the Triangle Normal Lemma, $\sin \angle(N_\tau, N_w\Sigma) \leq \frac{R}{\operatorname{lfs}(w)} \cot(45° - \phi/4) \leq \kappa \cot(45° - \phi/4)$. Then $\angle(N_\tau, N_{\tilde{x}}\Sigma) \leq \angle(N_{\tilde{x}}\Sigma, N_w\Sigma) + \angle(N_\tau, N_w\Sigma) \leq \eta(2\kappa) + \arcsin(\kappa \cot(45° - \phi/4))$. □

Chapter 3

Restricted Constrained Delaunay Triangulations

3.1 Introduction

The constrained Delaunay triangulation (CDT) in the plane [62, 23, 79] is a popular geometric construction that shares some of the advantages and mathematical properties of the Delaunay triangulation, but also permits users to constrain specified edges to be part of the triangulation. CDTs are used in applications such as computer graphics, geographical information systems, and guaranteed-quality mesh generation algorithms [20]. Our goal here is to offer a mathematically rigorous way to define a Delaunay-like triangulation on a curved surface embedded in three-dimensional space, with the same ability to constrain edges.

Another variant of the Delaunay triangulation, called the *restricted Delaunay triangulation* (RDT), has become a well-established way of generating triangulations on curved surfaces [36]. RDTs have equipped theorists to rigorously prove the correctness of algorithms for surface reconstruction [28] and surface mesh generation [20]. Here we introduce *restricted constrained Delaunay triangulations* (restricted CDTs), which combine ideas from CDTs and RDTs to enable the enforcement of constraining edges in RDTs. With restricted CDTs, it will be possible to improve algorithms for guaranteed-quality mesh generation on surface patches.

Think of the restricted CDT as a function that takes in three inputs: a compact, smooth surface $\Sigma \subset \mathbb{R}^3$ without boundary; a finite set $V \subset \Sigma$ of sample points (called *sites* or *vertices*); and a finite set S of line segments whose endpoints are in V. If certain conditions on the density of V and the lengths of the segments are met then, as illustrated in Figure 3.1, the output is a simplicial complex $\mathcal{T}$ such that the set of vertices of $\mathcal{T}$ is V, the set of edges of $\mathcal{T}$ is a superset of S, and $\mathcal{T}$ is a triangulation of Σ. By the last phrase, we mean that the *underlying space* of $\mathcal{T}$, written $|\mathcal{T}| = \bigcup_{\tau \in \mathcal{T}} \tau$, is homeomorphic to Σ.

Although Delaunay triangulations in the plane can be constrained to include arbitrary edges, the same is not true of three-dimensional Delaunay triangulations; consider the fact that not all nonconvex polyhedra can be tetrahedralized [77]. Nor is it always possible to constrain arbitrary

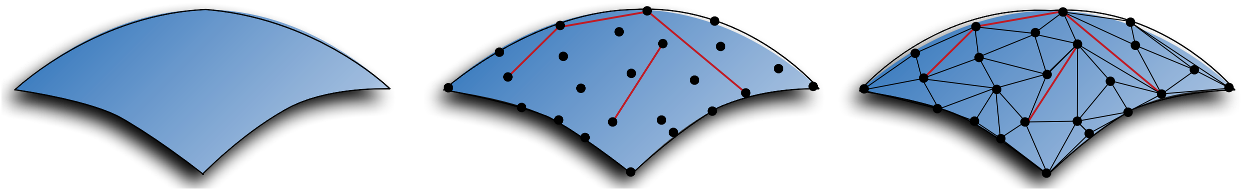

Figure 3.1: Given a set of points sampled from Σ and a set of segments, red, we wish to compute a triangulation of Σ that contains all of the red segments.

edges to be part of a surface triangulation. Our challenge is to establish conditions on the input that guarantee that a suitable triangulation exists.

We follow the example of the RDT, which is defined by dualizing a *restricted Voronoi diagram*. Given inputs Σ and V (but no segments), the *restricted Voronoi cell* of a site $v \in V$, denoted $\text{Vor}\,|_\Sigma v$, is the set of all points on Σ for which v is the closest site in V (possibly tied for closest), as measured by the Euclidean distance in $\mathbb{R}^3$. Equivalently, $\text{Vor}\,|_\Sigma v = \text{Vor}\, v \cap \Sigma$, where $\text{Vor}\, v$ is v's standard Voronoi cell in $\mathbb{R}^3$. The name "restricted Voronoi cell" arises because $\text{Vor}\,|_\Sigma v$ is the restriction of $\text{Vor}\, v$ to the surface Σ.

A *restricted Voronoi face* is any nonempty set of points found by taking the intersection of one or more restricted Voronoi cells. A *restricted Voronoi edge* (which is a curve on Σ) or a *restricted Voronoi vertex* (a one-point face) usually "belongs to" two or three equidistant closest sites in V. The *restricted Voronoi diagram* $\text{Vor}\,|_\Sigma V$ is the cell complex containing all the restricted Voronoi cells and faces.

The *restricted Delaunay triangulation* $\text{Del}\,|_\Sigma V$ is the simplicial complex dual to $\text{Vor}\,|_\Sigma V$. That is, a simplex is in $\text{Del}\,|_\Sigma V$ if and only if its vertices' restricted Voronoi cells have a nonempty mutual intersection. Roughly speaking, the RDT is found by drawing a triangle for each restricted Voronoi vertex on Σ; the triangle connects the three sites whose restricted Voronoi cells touch the restricted Voronoi vertex.

To modify RDTs so that we can constrain edges, we borrow from Seidel [79] the idea of an *extended Voronoi diagram*, which is the natural dual of the CDT. Seidel performs a topological surgery on the plane in which each segment in S becomes a slit cut in the plane; upon these slits he glues topological extensions called "secondary sheets" on which additional portions of the extended Voronoi diagram are drawn. Likewise, we perform surgery by cutting slits in the surface Σ and grafting independent new surfaces called *extrusions* onto Σ at these slits. We think of these slits as *portals*: an ant crawling on the surface across a constraining segment finds itself transported by the portal to an alternative space where the extended surface continues along an infinite extrusion. These extrusions are not merely topological; their geometry also requires careful definition.

Our main contribution is the definition of the restricted constrained Delaunay triangulation, which we define as the dual of the Voronoi diagram restricted to this surgically modified surface. We prove several combinatorial properties of restricted CDTs, including conditions under which the restricted CDT contains every constraining segment, conditions under which the restricted CDT is homeomorphic to the underlying surface Σ, and a characterization of which vertices must be considered to compute the triangles near a segment. The restricted CDT has immediate practical applications in surface meshing and geometric modeling, such as meshing trimmed splines.

An alternative approach sometimes suggested is to define a Voronoi diagram based on an intrinsic (geodesic) distance metric, then obtain a triangulation by duality. While this idea is mathematically elegant, it is not practical; computing a geodesic Voronoi diagram requires complicated and slow numerical discretization algorithms. That is why RDTs have been preferred for practical mesh generation applications: they are much easier to compute. We emphasize that although our mathematical construction of restricted CDTs may seem complicated, it is in the service of producing simple algorithms. In particular, computing the restricted Delaunay triangles near a segment on one "side" of the segment (specifically, triangles whose dual Voronoi vertices lie on a particular extrusion) is usually just a simple matter of ignoring the sites on the other "side." We do not have space to discuss algorithms here, but we think that the most practical algorithm for constructing a restricted CDT will first construct the RDT of V and Σ, then incrementally insert the segments of S one by one.

3.2 Portals and Topological Surgery

Informally, a *portal* P is a subset of a topological space X that acts as a doorway between subsets of X. We define X by starting with disjoint subsets Y and Z and then gluing them together by specifying an equivalence relationship between a subset of points $P \subset Y$ and a subset $P' \subset Z$. For clarity, we explain Seidel's construction of portals in the plane [79] before explaining our construction of portals on surfaces.

Portals in the Plane

Let $X = \mathbb{R}^2$ and let S be a finite set of line segments in the plane; the segments may intersect each other only at their endpoints. Consider a segment $s \in S$. The relative interior of s, which by a slight abuse of notation we denote by $\text{Int}\, s$, consists of all points of s except its two endpoints. Let the *slitted plane* $X_s = X - \text{Int}\, s$ be the plane with the interior of s removed. The affine hull of s, denoted $\text{aff}\, s$, has two "sides." Our goal is to augment X_s by gluing it to two additional topological spaces, one for each side of $\text{aff}\, s$, along the slit created by removing $\text{Int}\, s$. The three spaces are glued together along two portals, each of which is topologically a copy of s. Thus an ant crawling on the augmented space that enters s from one side finds itself in a *secondary branch*; and an ant that enters s from the other side finds itself in a different secondary branch. After repeating this augmentation for every segment in S, we can draw on the augmented space an *extended Voronoi diagram* whose dual is the CDT.

Let p and q be the endpoints of s. Topologically, X_s has a hole such that X_s is *almost* an open set, except that X_s has two boundary points, p and q. We want to glue two additional spaces to X_s—one for each side of s—so we augment X_s with additional points that serve as two portals to those additional spaces. We define a closed topological space $\overline{X}_s$ by augmenting X_s with two connected curves ζ_+ and ζ_-, called *portals*, that together serve as the boundary of the hole. Each of ζ_+ and ζ_- has p and q as its endpoints, but the two curves share no other points. In essence, the portals are copies of s with shared endpoints. $\overline{X}_s$ is the completion of the incomplete metric space X_s with respect to the shortest-path metric in X_s.

The points in X_s inherit Cartesian coordinates from the plane, and the points on the portals ζ_+ and ζ_- inherit Cartesian coordinates from the segment s. Two points in $\overline{X}_s$—one on ζ_+ and one on ζ_-—can have the same (x, y)-coordinate values yet be topologically distinct.

Let $\mathbb{R}^2_-$ and $\mathbb{R}^2_+$ be two copies of $\mathbb{R}^2$. We treat $\overline{X}_s$, $\mathbb{R}^2_-$, and $\mathbb{R}^2_+$ as three distinct topological spaces that all inherit the Cartesian coordinate system—so two points in two different spaces can have the same coordinate values yet be topologically distinct.

Informally, we glue $\mathbb{R}^2_+$ to $\overline{X}_s$ along ζ_+ and glue $\mathbb{R}^2_-$ to $\overline{X}_s$ along ζ_-. Formally, we write $x \equiv y$ if x and y have the same coordinate values, even though they may lie in different spaces. Let p and q be the endpoints of s. Define an equivalence relation $\sim$ as

$$x \sim y \iff \begin{cases} x = y & x, y \in \overline{X}_s \text{ or } x, y \in \mathbb{R}^2_+ \text{ or } x, y \in \mathbb{R}^2_-, \\ x \equiv y & x \in \mathbb{R}^2_+ \text{ and } y \in \zeta_+, \\ x \equiv y & x \in \mathbb{R}^2_- \text{ and } y \in \zeta_-, \\ x \equiv p \equiv y \text{ or } x \equiv q \equiv y & x \in \mathbb{R}^2_+ \text{ and } y \in \mathbb{R}^2_-. \end{cases}$$

With $\sim$ we construct the quotient space $\widetilde{X} = (\overline{X}_s \sqcup \mathbb{R}^2_+ \sqcup \mathbb{R}^2_-)/\sim$. We refer to $\overline{X}_s$ as the *principal branch* and refer to $\mathbb{R}^2_+$ and $\mathbb{R}^2_-$ as *secondary branches*. Figures 3.2 and 3.3 illustrate this construction. Note that the endpoints of the segment s in the quotient space are present in, and shared by, all three of the original spaces.

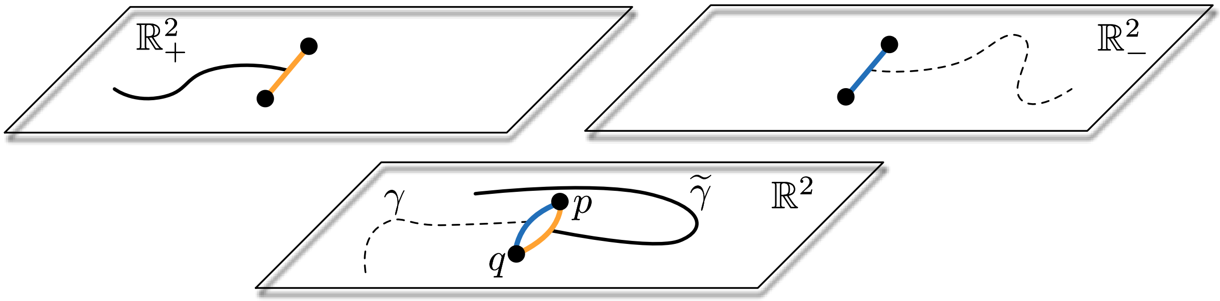

Figure 3.2: The completion of the slitted plane has a topological hole bounded by two portals, marked in blue and orange. (Geometrically, the two portals are straight line segments that occupy exactly the same coordinates.) The equivalence relation $\sim$ identifies the blue path in the principal branch with the blue path in $\mathbb{R}^2_-$; likewise the two orange paths become one. A path in the principal branch (bottom) that enters a portal continues in the appropriate secondary branch.

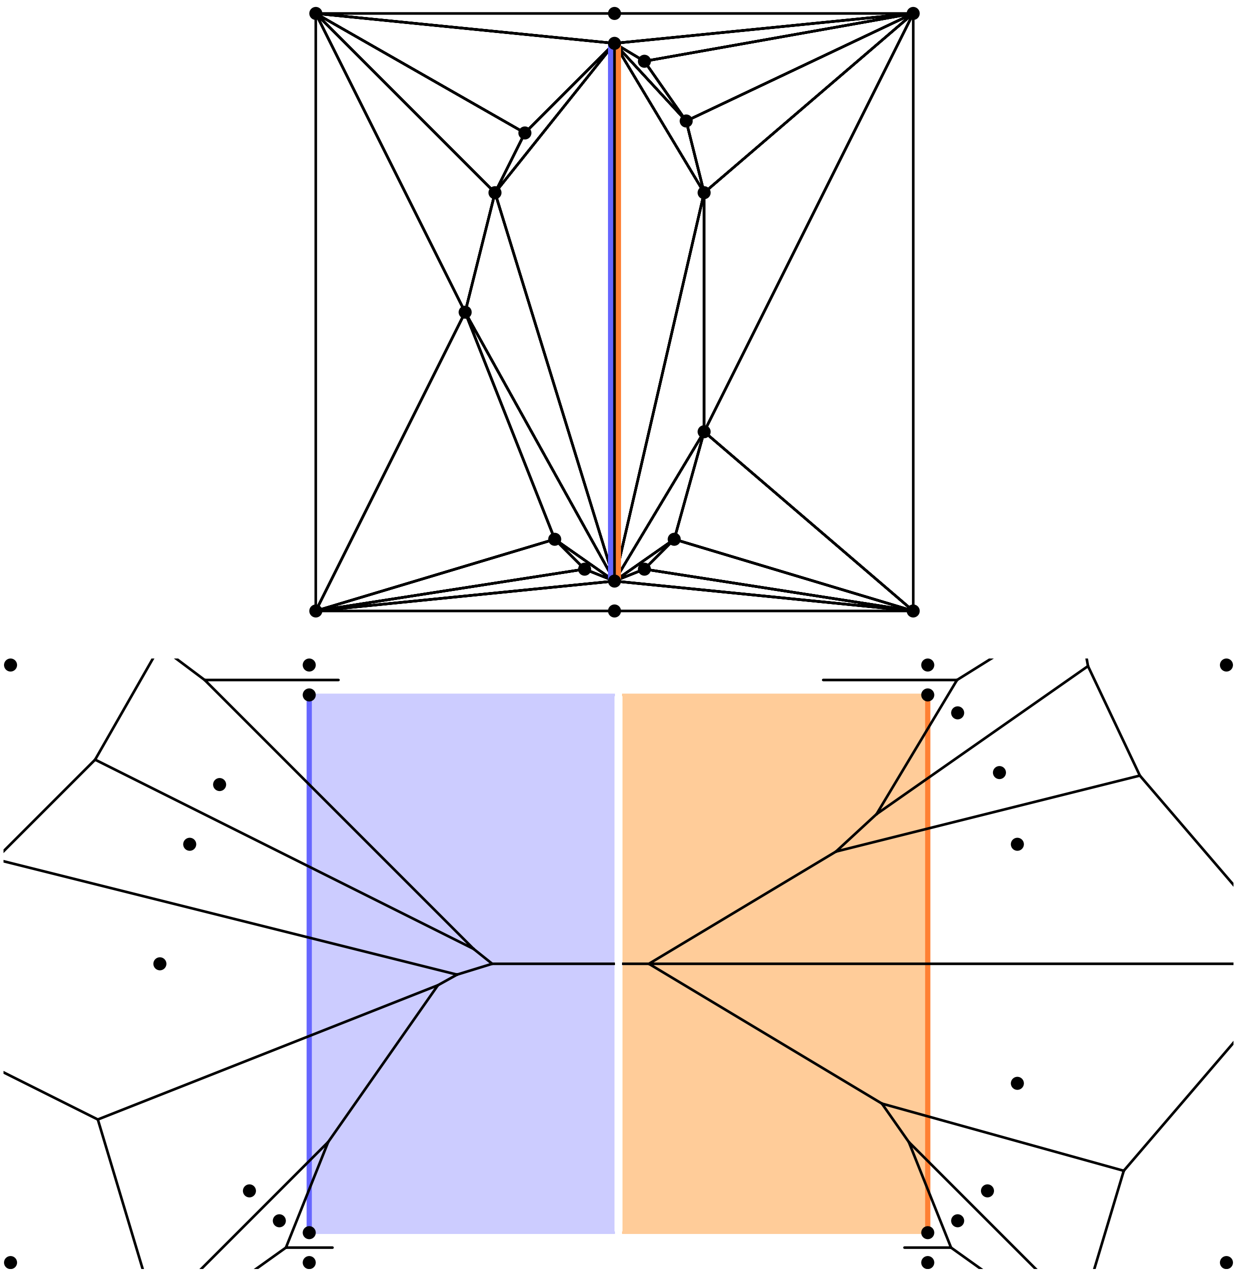

Figure 3.3: A one-segment CDT and its dual extended Voronoi diagram. The blue and orange regions show the portions of the Voronoi diagram on the secondary branches.

The construction works for any finite number $m = |S|$ of non-crossing segments. We remove the segment interiors from X, $X_S = X - \bigcup_{s \in S} \operatorname{Int} s$, then we add two portals for each segment (taking the completion of X_S) to yield $\overline{X}_S$. Then we construct a quotient space $\widetilde{X}$ composed of $\overline{X}_S$ and $2m$

copies of $\mathbb{R}^2$ glued along the $2m$ portals bounding the m holes in $\overline{X}_S$.

For the sake of defining the Voronoi diagram of a finite set of sites in $\widetilde{X}$, Seidel [79] defines a distance function on $\widetilde{X}$ which is essentially the Euclidean distance, except that the distance between two points is infinite if they are not *visible* from each other. (Note that this distance function is not a metric.) Consider a continuous, injective curve $\gamma : [0, 1] \to \widetilde{X}$. The curve γ may pass through portals and visit secondary branches, but because of the slits we have cut in X_S, γ cannot cross the interior of a segment without being transported by a portal. We call a curve *straight* if its Cartesian embedding is a straight line segment. Two points $p, q \in \widetilde{X}$ are visible from each other if there is a straight curve $\gamma \subset \widetilde{X}$ with endpoints p and q. The distance $\widehat{d}(p, q)$ from p to q is the Euclidean distance if p and q are visible from each other; otherwise, $\widehat{d}(p, q) = \infty$.

The extended Voronoi diagram assigns each point in $\widetilde{X}$ to (the Voronoi cells of) one or more sites in V. Those sites must be visible from the point; no site can claim a point it cannot see. If a point on a secondary branch is claimed by a site other than the branch's portal's endpoints, the site must be visible from the point through the portal. Seidel gives an algorithm for constructing the extended Voronoi diagram, and by duality the CDT.

Portals on Surfaces Embedded in $\mathbb{R}^3$

A similar construction works for a smooth surface $\Sigma \subset \mathbb{R}^3$. However, whereas in the plane we build one new topological space, here we will require two. We surgically augment the surface Σ by cutting slits along *portal curves*, one for each segment, and gluing two *extrusions* onto each portal curve, yielding an extended surface $\widetilde{\Sigma}$. The purpose of this extended surface is to serve as a canvas upon which we can draw an extended restricted Voronoi diagram, which we can dualize to define a restricted CDT of Σ and S.

However, to define a Voronoi diagram we need a distance function, and $\widetilde{\Sigma}$ alone does not suffice to define a useful distance function. Recall that while an intrinsic (geodesic) distance might be ideal in principle, it is not practical for fast computation. For the sake of speed and practicality, we wish to use the Euclidean distance in $\mathbb{R}^3$ as RDTs do, but the Euclidean distance must be modified so that the restricted Voronoi diagram respects the input segments. Hence most of our work will be to construct a surgically modified three-dimensional space $\widetilde{X}$ in which we embed $\widetilde{\Sigma} \subset \widetilde{X}$. Like Seidel's augmented space in Section 3.2, $\widetilde{X}$ obstructs (and supports) visibility in a manner suitable for defining a restricted Voronoi diagram on $\widetilde{\Sigma}$, and does so in a way that makes it easy to compute restricted CDTs.

To define $\widetilde{X}$, we specify portals in $\mathbb{R}^3$ where points will be removed, analogous to cutting slits in the plane. Each portal is a two-dimensional ruled surface with boundary (not generally flat) which is a union of line segments. Each portal curve is the intersection of a portal with Σ. Each line segment is perpendicular to Σ where it intersects Σ, on a portal curve. Each portal has two "sides," and on each side we glue an additional copy of $\mathbb{R}^3$ in which we embed an extrusion. The extended Voronoi diagram assigns each point in $\widetilde{\Sigma}$ to one or more sites in V that are visible from the point along a straight path in $\widetilde{X}$.

Let $\Sigma \subset \mathbb{R}^3$ be a compact, smooth surface without boundary. The *medial axis* M of Σ is the closure of the set of all points in $\mathbb{R}^3$ for which the closest point on Σ is not unique. Intuitively, the medial axis of Σ is meant to capture the "middle" of the region bounded by Σ. A *medial ball* is a ball whose center lies on M and whose boundary intersects Σ (tangentially), but the interior of the ball does not. For any point $x \in \Sigma$, there are one or two medial balls that have x on their boundaries, called *medial balls at* x. If there are two, there is one on each side of Σ.

For $x \in \Sigma$, the *normal segment* ℓ_x at x is a line segment or ray that passes through x, is orthogonal to Σ at x, and has its endpoints on M. If there are two medial balls at x, the endpoints of ℓ_x are the centers of those medial balls. If there is only one medial ball at x, then ℓ_x is a ray originating at the medial ball's center.

The *local feature size* function is $\text{lfs} : \Sigma \to \mathbb{R}, x \mapsto d(x, M)$ where $d(x, M)$ denotes the Euclidean distance from x to M. We require that Σ is smooth in the sense that $\inf_{x\in\Sigma} \text{lfs}(x) > 0$. A finite point set $V \subset \Sigma$ is an *ϵ-sample of* Σ if for every point $x \in \Sigma$, $d(x, V) \leq \epsilon\, \text{lfs}(x)$. That is, the ball with center x and radius $\epsilon\, \text{lfs}(x)$ contains at least one sample point. See Figure 3.4.

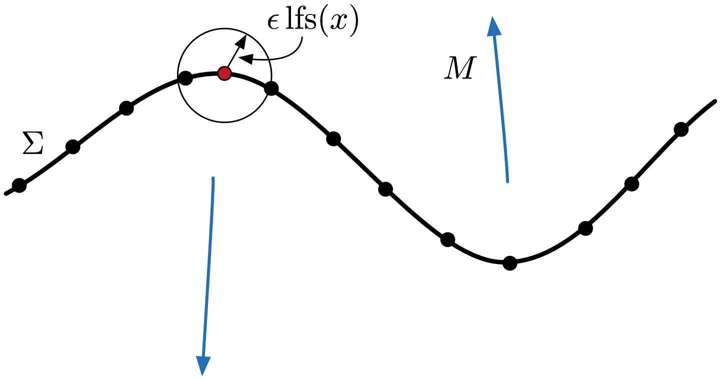

Figure 3.4: The medial axis M (blue) of a curve Σ embedded in the plane (as three-dimensional examples are hard to draw or understand). The medial axis is unbounded; both depicted components extend infinitely far away. The set of black vertices is a 0.5-sample of Σ. The point $x \in \Sigma$, shown in red, has at least one sample point in the ball centered at x with radius $\epsilon\, \text{lfs}(x)$.

Let S be a finite set of line segments whose endpoints lie on Σ. Let $s \in S$ be a segment with endpoints $p, q \in \Sigma$. Let B_s be the *diametric ball* of s—the smallest closed ball such that $s \subset B_s$, so that s is a diameter of B_s. Suppose that $d(p, q) \leq \rho\, \text{lfs}(p)$ for some $\rho \in (0, 1)$; that is, s is short relative to the local feature size. Then one can prove that $B_s \cap \Sigma$ is a topological disk (see Lemma 40 in Appendix A.1).

Suppose that we know or can approximate the unit vector n_p normal to Σ at any site p. We choose a *cutting plane* $h_s \supset s$ that is locally orthogonal to the surface Σ at p or q (or perhaps somewhere between p and q). We use h_s to specify a *portal curve* $\zeta_s = h_s \cap B_s \cap \Sigma$, which is a single connected curve from p to q on Σ. There is not a canonical choice of cutting plane (and thus portal curve) for s, and the user might be presented with a range of choices, but for our presentation here, we choose $h_s = \text{Span}\{n_p, \vec{pq}\}$. We require that the portal curves do not cross each other. More precisely, the relative interior of a portal curve may not intersect another portal curve nor a site in V.

Our requirement that each portal curve must lie on a plane has both a theoretical motivation and a practical one. Our proof that every constraining segment is an edge in the restricted CDT (Theorem 19) depends on the fact that each portal curve lies in a plane and its extrusions are orthogonal to that plane. The requirement simplifies algorithms for computing a restricted CDT, because the Voronoi cells on an extrusion are solely influenced by sites on the other side of the cutting plane—plus p and q. (See Theorem 18.)

Figure 3.5 illustrates our portal construction. For each segment s, the portal $P_s = \bigcup_{x \in \zeta_s} \ell_x$ is the union of the normal segments of the points on the portal curve ζ_s. Hence a portal is a ruled surface, topologically two-dimensional but not lying in a plane. Each portal reaches to the medial axis, thereby obstructing visibility so that sites on one "side" of a segment do not affect the restricted Delaunay triangles on the other "side."

We construct the augmented space $\widetilde{X}$ as we did in Section 3.2, with P_s replacing s and $\mathbb{R}^3$ replacing $\mathbb{R}^2$. We first let $X = \mathbb{R}^3$. We construct the space $X_s = X - \text{Int}\, P_s$, which is $\mathbb{R}^3$ with the relative interior of P_s removed. We construct $\overline{X}_s$ (the completion of X_s) by augmenting the slit with two portals, one for each side of P_s. Each portal is a topological copy of P_s, but both copies share a single boundary ∂P_s with each other and with X_s. Then we construct $\widetilde{X}$ by gluing two secondary branches (copies of $\mathbb{R}^3$) to X_s, one along each portal.

The construction works for any finite number of segments, so long as no portal curve intersects the relative interior of another portal curve, which implies that no portal intersects the relative interior of another portal. If two segments share an endpoint p, then their portals share a boundary segment ℓ_p; and two portals' boundaries may also intersect each other at the medial axis; and these boundary points are considered equivalent in the quotient space. However, no portal interiors intersect each other.

For a finite number of segments, the construction removes the slits for all segments from X, $X_S = X - \bigcup_{s \in S} \text{Int}\, P_s$, then creates two portals for each segment to yield the completion $\overline{X}_S$. Then we construct the quotient space $\widetilde{X}$ with $2m$ copies of $\mathbb{R}^3$ glued along the $2m$ portals bounding the m holes in $\overline{X}_S$.

Our next step is to surgically modify Σ to create an extended surface $\widetilde{\Sigma}$ embedded in $\widetilde{X}$, as illustrated at bottom in Figure 3.5. Let b_s be a unit vector normal to the cutting plane h_s. We extrude the portal curve ζ_s into each of the two secondary branches connected to the portal P_s, in the two directions normal to h_s. For each point $x \in \zeta_s$ we extrude a ray in the direction of b_s into $\mathbb{R}^3_+$, and another in the direction of $-b_s$ into $\mathbb{R}^3_-$. More precisely, we define the ruled surfaces $\Sigma^+_s = \{x + \omega b_s \in \mathbb{R}^3_+ : x \in \zeta_s, \omega \in [0, \infty)\}$ and $\Sigma^-_s = \{x - \omega b_s \in \mathbb{R}^3_- : x \in \zeta_s, \omega \in [0, \infty)\}$. Let $\Sigma_S = \Sigma - \bigcup_{s \in S} \text{Int}(\zeta_s)$ be the surface with the portal curve interiors removed, and let $\overline{\Sigma}_S$ be its completion. Define an equivalence relation $\sim$ that identifies points along ζ_s on $\overline{\Sigma}_S$, each Σ^+_s, and each Σ^-_s. Our extended surface is $\widetilde{\Sigma} = \left(\overline{\Sigma}_S \sqcup \bigsqcup_{s \in S} \Sigma^+_s \sqcup \bigsqcup_{s \in S} \Sigma^-_s\right) / \sim$.

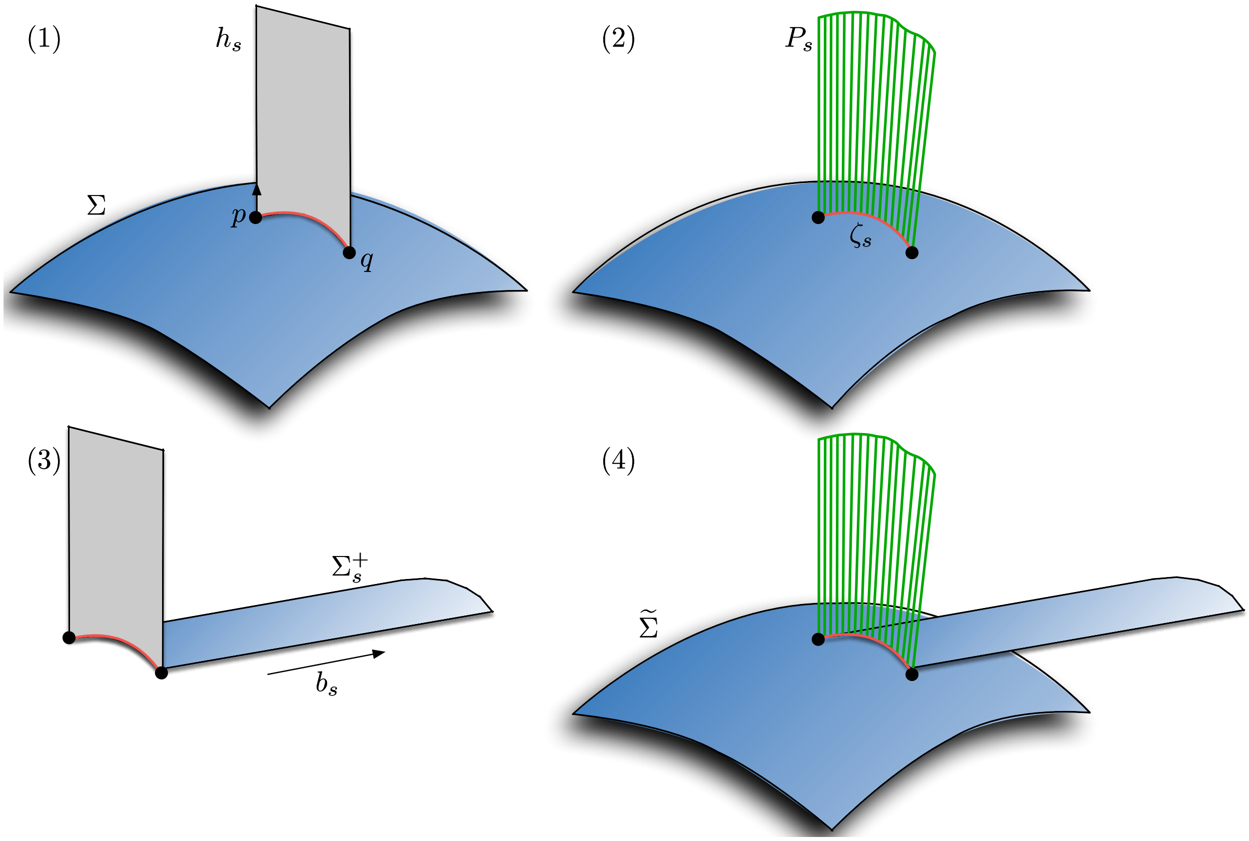

Figure 3.5: (1) The plane h_s intersects Σ in a curve; the portal curve ζ_s (red) is the portion of this curve in the diametric ball B_s of the segment s. (2) Our portal P_s, shown in green, is the union of the normal segments (locally orthogonal to Σ) of the points on the portal curve ζ_s. The normal segments terminate on the medial axis M. (3) We extrude the portal curve ζ_s into $\mathbb{R}^3_+$ in the direction b_s orthogonal to h_s, thus defining Σ_s^+. (4) We glue the extrusion Σ_s^+ to Σ along ζ_s at the entrance to the portal P_s.

3.3 The Geometry of Portal Curves

In this section we derive some facts about the geometry of portal curves that we use to prove this chapter's main results, including Theorems 18, 20, and 23.

Some Facts About Plane Geometry and Curvature

Let s be a line segment in the plane with endpoints p and q, let C be the circle whose diameter is s, let c be the center of C (the midpoint of s), and let v be a vector normal to s. The bisector of s is the line $l = \{c + \lambda v : \lambda \in \mathbb{R}\}$. Let $\mathcal{P}(s)$ denote the pencil of circles that pass through both p and q. All these circles have their centers on l.

Without loss of generality, we choose a coordinate system so that c (the midpoint of s/center of C) is the origin and s lies on the x-axis. C is the unique smallest circle in the pencil $\mathcal{P}(s)$, and its radius is $d(p, q)/2$. For every value $r > d(p, q)/2$, there are two circles $C_r, C_{-r} \in \mathcal{P}(s)$ with radius r; here we define C_r to have its center above s (i.e., with positive y-coordinate) and C_{-r} to have its center below s, as illustrated in Figure 3.6.

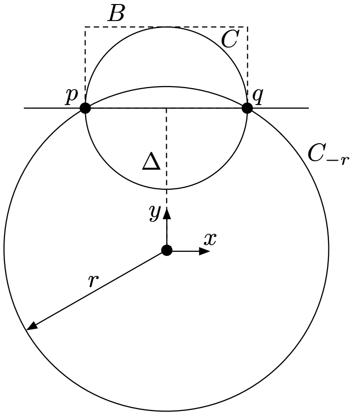

Figure 3.6: The circles C and C_{-r} are members of the pencil of circles $\mathcal{P}(pq)$ induced by the segment pq. We define a coordinate system whose origin is the center of C_{-r} and whose x-axis is parallel to pq.

Lemma 9. *Let $p, q \in \mathbb{R}^2$ be two points and let C be the circle with diameter pq. Let $C_{-1/\kappa} \in \mathcal{P}(pq)$ be a circle with radius $1/\kappa$ (curvature κ). Then the distance from the north pole of $C_{-1/\kappa}$ to the line segment pq is*

$$\frac{1 - \sqrt{1 - \kappa^2\, d(p,q)^2/4}}{\kappa}.$$

By symmetry, the same bound holds for the south pole of $C_{1/\kappa}$.

Proof. Let $c = (0, \Delta)$ be the center of $C_{-1/\kappa}$. As $C_{-1/\kappa}$ passes through p and q, its radius is $1/\kappa = d(c, p) = d(c, q) = \sqrt{d(p, q)^2/4 + \Delta^2}$. Solving for Δ gives

$$\Delta = \pm\frac{\sqrt{1 - \kappa^2\, d(p,q)^2/4}}{\kappa}.$$

The negative solution is the relevant one. (The positive solution provides the symmetric result for the south pole of $C_{1/\kappa}$.) The distance between the north pole of $C_{-1/\kappa}$ and the line segment pq is

$$\left|\Delta + \frac{1}{\kappa} - 0\right| = \left|\frac{1 - \sqrt{1 - \kappa^2\, d(p,q)^2/4}}{\kappa}\right| = \frac{1 - \sqrt{1 - \kappa^2\, d(p,q)^2/4}}{\kappa}.$$

□

Given a circle C, let $B(C)$ denote the closed disk whose boundary is C. The following lemma states that if the curvature of a portal curve ζ_s is bounded, then $\zeta_s \subset B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$. Intuitively, ζ_s cannot stray far from the line segment $s = pq$. This result allows us to bound the distance between ζ_s and s on a portal P_s.

Lemma 10. *Let $p, q \in \mathbb{R}^2$ be two points and let C be the circle with diameter pq. Let $\gamma : [0, 1] \to \mathbb{R}^2$ be a regular curve in $\mathbb{R}^2$ with $\gamma(0) = p$ and $\gamma(1) = q$ and with curvature everywhere at most κ. Suppose that $\gamma \subset B(C)$. Then $\gamma \subset B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$.*

Proof. First we show that $\gamma \subset B(C_{-1/\kappa})$; suppose for the sake of contradiction that $\gamma \not\subset B(C_{-1/\kappa})$. Let $\gamma(t')$ be a point on γ that maximizes the distance from $\gamma(t')$ to the center of $C_{-1/\kappa}$; then $\gamma(t') \notin B(C_{-1/\kappa})$. Let C' be a circle of radius $1/\kappa$ tangent to γ at $\gamma(t')$ with the center of C' lying on the line segment connecting $\gamma(t')$ to the center of $C_{-1/\kappa}$. As C' has the same radius as $C_{-1/\kappa}$ (which is greater than the radius of C) but passes through a point in $C \setminus C_{-1/\kappa}$, the circle C' must enclose either p or q (possibly both). But γ is a curve in C whose curvature nowhere exceeds κ, and it is tangent to C' at a point, so no part of γ can be inside C'. By contradiction, it follows that $\gamma \subset B(C_{-1/\kappa})$. By a symmetrical argument, $\gamma \subset B(C_{1/\kappa})$. □

For any point in $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$, we can bound the distance between that point and the nearest endpoint of pq. The bound is at its worst on the boundary of $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$ so we compute the maximum distance to one of the endpoints for every point along the relevant arcs of $C_{1/\kappa}, C_{-1/\kappa}$.

For the next lemma, we find it convenient to define a coordinate system different from the one used in the previous two lemmas. Consider a coordinate system where the center of C_{-r} is the origin, the center of C lies on the positive y-axis, and the segment pq is parallel with the x-axis, as in Figure 3.6. In this coordinate system, $p = (-d(p, q)/2, \Delta)$ and $q = (d(p, q)/2, \Delta)$.

Lemma 11. *Let $p, q \in \mathbb{R}^2$ be two points and let C denote the circle with diameter pq. Parameterize C as*

$$C(\theta) = \left(-\frac{d(p, q)}{2} \cos\theta, \Delta + \frac{d(p, q)}{2} \sin\theta\right).$$

For each value of $\theta \in [0, \pi]$, the line segment from the center of C to $C(\theta)$ intersects a unique point on C_{-r}. The distance from this point on C_{-r} to p or q (whichever is closest) can be expressed in *terms of θ as*

$$\min\{d(C_{-r}(\theta),p),d(C_{-r}(\theta),q)\} = \begin{cases} [r^2 + \frac{d(p,q)^2}{4} - \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2} \\ \quad - \cos 2\theta(\Delta^2 + \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2}) \\ \quad - \Delta\sin 2\theta(-\frac{d(p,q)}{2} + \sqrt{r^2\sec^2\theta - \Delta^2}]^{1/2} & \theta \in [0,\pi/2), \\ \sqrt{r^2 + \Delta^2 - 2r\Delta + \frac{d(p,q)^2}{4}} & \theta = \pi/2, \\ [r^2 + \frac{d(p,q)^2}{4} - \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2} \\ \quad - \cos 2\theta(\Delta^2 + \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2}) \\ \quad + \Delta\sin 2\theta(-\frac{d(p,q)}{2} + \sqrt{r^2\sec^2\theta - \Delta^2})]^{1/2} & \theta \in (\pi/2,\pi]. \end{cases} \tag{3.1}$$

A symmetric result holds for C_r.

Proof. C_{-r} can be parameterized as

$$C_{-r}(\varphi) = (-r\cos\varphi, r\sin\varphi).$$

Then the squared distance from a point on C_{-r} to p is

$$d(C_{-r}(\varphi),p)^2 = r^2 + \Delta^2 + \frac{d(p,q)^2}{4} - rd(p,q)\cos\varphi - 2r\Delta\sin\varphi.$$

Similarly, the squared distance from a point on C_{-r} to q is

$$d(C_{-r}(\varphi),q)^2 = r^2 + \Delta^2 + \frac{d(p,q)^2}{4} + rd(p,q)\cos\varphi - 2r\Delta\sin\varphi.$$

We wish to reparameterize these two distances in terms of θ; that is, for each $\theta \in [0,\pi]$ we need an expression for the corresponding value of φ. It follow from basic trigonometry that

$$\tan\theta = \frac{|r\sin\varphi - \Delta|}{|r\cos\varphi|}. \tag{3.2}$$

We'll consider the solutions to this equation in two steps; first for $\varphi \in [\arcsin(\Delta/r),\pi/2)$ and second for $\varphi \in (\pi/2, \pi - \arcsin(\Delta/r)]$. In both of these intervals, $r\sin\varphi - \Delta \geq 0$, so we can remove the absolute value from the numerator in Equation 3.2. In the case where $\varphi \in [\arcsin(\Delta/r),\pi/2)$, $r\cos\varphi > 0$ so we can also remove the absolute values in the denominator. In this case,

$$\varphi = \arctan\frac{\Delta + \tan\theta\sqrt{r^2\sec^2\theta - \Delta^2}}{-\Delta\tan\theta + \sqrt{r^2\sec^2\theta - \Delta^2}}.$$

In the second case, $\varphi \in (\pi/2, \pi - \arcsin\frac{\Delta}{r}]$, $r\cos\varphi < 0$. Thus we replace the absolute value in the denominator with a negative sign, and

$$\varphi = \arctan\frac{\Delta - \tan\theta\sqrt{r^2\sec^2\theta - \Delta^2}}{\Delta\tan\theta + \sqrt{r^4\sec^2\theta - \Delta^2}}.$$

However, we're interested in the angle that the hypotenuse makes with the negative x-axis. Thus the correct value of φ is

$$\varphi = \pi - \arctan \frac{\Delta - \tan\theta \sqrt{r^2 \sec^2\theta - \Delta^2}}{\Delta \tan\theta + \sqrt{r^4 \sec^2\theta - \Delta^2}}.$$

Plugging our two solutions for φ into our two distance equations respectively, we have

$$\begin{aligned} r^2 + \frac{d(p,q)^2}{4} - \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2} \\ - \cos 2\theta\left(\Delta^2 + \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2}\right) \\ - \Delta\sin 2\theta\left(-\frac{d(p,q)}{2} + \sqrt{r^2\sec^2\theta - \Delta^2}\right) \end{aligned}$$

and

$$\begin{aligned} r^2 + \frac{d(p,q)^2}{4} - \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2} \\ - \cos 2\theta\left(\Delta^2 + \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2}\right) \\ + \Delta\sin 2\theta\left(-\frac{d(p,q)}{2} + \sqrt{r^2\sec^2\theta - \Delta^2}\right). \end{aligned}$$

The first equation holds in the range $[0, \pi/2)$, while the second holds in the range $(\pi/2, \pi]$. Taking square roots gives Equation 3.1.

The point corresponding to $\theta = \varphi = \pi/2$ is equidistant from both p and q. Its value can be computed by considering the left and right hand limits and confirming that they are identical. Indeed the limits are equal and have value

$$d(C_{-r}(\theta), p) = d(C_{-r}(\theta), q) = \sqrt{r^2 + \Delta^2 - 2r\Delta + \frac{d(p,q)^2}{4}}.$$

□

Curvature Bounds

For a segment s and plane h_s, defined as in Section 3.2, we wish to bound the curvature κ of the curve ζ_s. Using the fact that ζ_s is defined as the intersection of h_s with Σ, we bound the curvature κ in terms of the maximum principal curvatures of Σ along ζ_s.

Lemma 12. *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface and let $p, q \in \Sigma$ be points on Σ. Let $h = \mathrm{Span}\{n_p, \vec{pq}\}$ be a plane and define $b_h = \frac{n_p \times \vec{pq}}{\|n_p \times \vec{pq}\|}$, the unit vector normal to h. Define a curve $\gamma : [0,1] \to \Sigma$ as the intersection $h \cap \Sigma$ such that $\gamma(0) = p$ and $\gamma(1) = q$. Let κ be the curvature along γ. Then the geodesic curvature is $\kappa_g = \kappa\langle b_h, n\rangle$.*

Proof. Let T denote the tangent field along γ and N denote the principal normal along γ. Then the geodesic curvature is

$$\begin{aligned}\kappa_g &= \kappa\langle N, n \times T\rangle \\ &= \kappa\langle b_h \times T, n \times T\rangle \\ &= \kappa \begin{vmatrix} \langle b_h, n\rangle & \langle b_h, T\rangle \\ \langle T, n\rangle & 1 \end{vmatrix} \\ &= \kappa(\langle b_h, n\rangle - \langle b_h, T\rangle\langle T, n\rangle).\end{aligned}$$

The third identity uses the scalar quadruple product. Notice that $\langle b_h, T\rangle = 0$, since T always lies in h, and $\langle T, n\rangle = 0$, since T is also a tangent field along a curve in Σ. The result follows. □

Lemma 13. *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface and let $p, q \in \Sigma$ be points on Σ. Let $h = \mathrm{Span}\{n_p, \vec{pq}\}$ be a plane and define $b_h = \frac{n_p \times \vec{pq}}{\|n_p \times \vec{pq}\|}$, the unit vector normal to h. Define a curve $\gamma : [0, 1] \to \Sigma$ as the intersection $h \cap \Sigma$ such that $\gamma(0) = p$ and $\gamma(1) = q$. Suppose that, for every $t \in [0, 1]$, the angle $\angle(n_p, n_{\gamma(t)}) \leq \alpha$ in radians. Then the curvature κ of the curve γ is bounded by*

$$\frac{\max\{\kappa_1, \kappa_2\}}{\sqrt{1 - \sin^2 \alpha}}$$

where κ_1, κ_2 are the principal curvatures of Σ.

Proof. The angle $\angle(n_p, n_{\gamma(t)}) \leq \alpha$ for all t, which implies that the the angle $\angle(b_h, n_{\gamma(t)}) \in \left[\frac{\pi}{2} - \alpha, \frac{\pi}{2} + \alpha\right]$. It follows that the cosine of the angle is in the range $[-\sin\alpha, \sin\alpha]$.

Define θ up to sign by $\langle b_h, n\rangle = \cos\theta$. Then $\kappa_g = \kappa\cos\theta$ and $\kappa_n = \kappa\sin\theta$ [70]. So

$$\begin{aligned}\kappa &= \frac{\kappa_n}{\sin\theta} \\ &= \frac{|\kappa_n|}{\sqrt{1 - (\langle b_h, n\rangle)^2}} \\ &\leq \frac{\max\{\kappa_1, \kappa_2\}}{\sqrt{1 - \sin^2\alpha}}.\end{aligned}$$

The inequality in the last line follows from the fact that the normal curvature is bounded by the maximum of the principal curvatures. □

Bounds in Terms of Local Feature Size

A normal plane Π at a point $x \in \Sigma$ is a plane containing the normal vector n_x. It follows that Π must necessarily contain a unique unit tangent vector of $T_x\Sigma$. The plane Π cuts Σ in a plane curve. The curvature of this plane curve is defined as the reciprocal of the radius of the osculating circle at x. Notice that the radius of the osculating circle is at least $\mathrm{lfs}(x)$ since there is an empty ball of radius $\mathrm{lfs}(x)$ tangent to Σ at x. This holds for any normal plane at x. In particular it holds for the planes of principal curvature. Thus the maximum principal curvature at x is at most $\frac{1}{\mathrm{lfs}(x)}$. Combining this observation with Lemma 13 we derive an upper bound on the curvature κ of the plane curve γ.

Lemma 14. *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface and let $p, q \in \Sigma$ be points on Σ such that $d(p, q) \leq \rho \operatorname{lfs}(p)$ for $\rho < 1/2$. Let $h = \operatorname{Span}\{n_p, \vec{pq}\}$ be a plane and define $b_h = \frac{n_p \times \vec{pq}}{\|n_p \times \vec{pq}\|}$, the unit vector normal to h. Define a curve $\zeta : [0, 1] \to \Sigma$ as the intersection $h \cap \Sigma$ such that $\zeta(0) = p$ and $\zeta(1) = q$ and ζ is included in the ball $B = B((p + q)/2, d(p, q)/2)$. Then the curvature κ of the curve ζ is at most*

$$\frac{1}{\operatorname{lfs}(p)\sqrt{1 - 2\rho}}.$$

Proof. For any point r on Σ, the medial balls at r impose a lower bound on the radius of any osculating circle generated by the intersection of Σ with a plane normal to Σ at r. In particular, the radii of the osculating circles that define the principal curvatures are also at least the local feature size. Applying the bound in Lemma 13, we have

$$\begin{aligned}\kappa &\leq \frac{\max\{\kappa_1, \kappa_2\}}{\sqrt{1 - \sin^2\alpha}} \\ &\leq \frac{1}{\min_{r\in\zeta} \operatorname{lfs}(r)} \frac{1}{\sqrt{1 - \sin^2\alpha}}.\end{aligned}$$

Suppose that $r^* = \arg\min_{r\in\zeta} \operatorname{lfs}(r)$, minimizes the local feature size along ζ. Since ζ is entirely included in the ball B, $d(p, r^*) \leq \rho \operatorname{lfs}(p)$. Then, by the Feature Translation Lemma (Lemma 38), we have that

$$\begin{aligned}\operatorname{lfs}(p) &\leq \frac{1}{1-\rho}\operatorname{lfs}(r^*) \\ (1-\rho)\operatorname{lfs}(p) &\leq \operatorname{lfs}(r^*).\end{aligned}$$

Thus we make the bound weaker by writing it as

$$\kappa \leq \frac{1}{(1-\rho)\operatorname{lfs}(p)} \frac{1}{\sqrt{1 - \sin^2\alpha}}.$$

Finally, by the Normal Variation Lemma (Lemma 4), $\angle(n_p, n_q) \leq \frac{\rho}{1-\rho}$. Furthermore, $\sin \frac{\rho}{1-\rho} \leq \frac{\rho}{1-\rho}$ for $\rho < 1$. (The values are essentially identical for $\rho < \frac{2}{5}$.) Our final bound is

$$\begin{aligned}\kappa &\leq \frac{1}{(1-\rho)\operatorname{lfs}(p)\sqrt{1 - \left(\frac{\rho}{1-\rho}\right)^2}} \\ &= \frac{1}{(1-\rho)\operatorname{lfs}(p)\sqrt{\frac{1-2\rho}{(1-\rho)^2}}} \\ &= \frac{1}{\operatorname{lfs}(p)\sqrt{1-2\rho}}.\end{aligned}$$

□

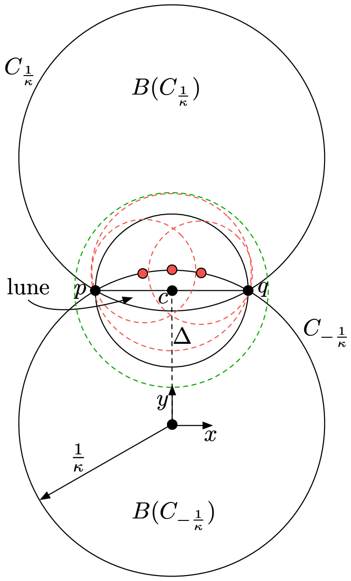

Figure 3.7: The circles $C_{1/\kappa}$ and $C_{-1/\kappa}$, and the disks $B(C_{1/\kappa})$ and $B(C_{-1/\kappa})$ bounded by these circles, all have radius $1/\kappa$. By Lemma 10, if the curvature of ζ_{pq} is at most κ, ζ_{pq} lies in the lune $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$. At each point z in the lune, we place a circle whose radius is $\min\{d(p,z), d(q,z)\}$. Three such circles are drawn in red. The union of all these circles is included in the green circle, centered at the midpoint of pq.

Recall Lemma 10, which states that γ is included in $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$, shown pictorially in Figure 3.7. We call this region the *lune*. One important feature of the lune is its *width*, the distance from the midpoint of pq to the north pole of $C_{-1/\kappa}$ (or equivalently to the south pole of $C_{1/\kappa}$). This distance was derived in terms of κ in Lemma 9. Using the bound in Lemma 14, we derive an upper bound on the width in terms of the local feature size.

Lemma 15. *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface and let $p, q \in \Sigma$ be points on Σ such that $d(p,q) \leq \rho\,\mathrm{lfs}(p)$ for $\rho < 1/2$. Let $h = \mathrm{Span}\{n_p, \vec{pq}\}$ be a plane and define $b_h = \frac{n_p \times \vec{pq}}{\|n_p \times \vec{pq}\|}$, the unit vector normal to h. Define a curve $\gamma : [0,1] \to \Sigma$ as the intersection $h \cap \Sigma$ such that $\gamma(0) = p$ and $\gamma(1) = q$ and γ is included in the ball $B = B((p+q)/2, d(p,q)/2)$. Then the width of the lune $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$ is at most*

$$\sqrt{1-2\rho}\left(1 - \sqrt{1 - \frac{\rho^2}{4(1-2\rho)}}\right)\mathrm{lfs}(p).$$

Proof. The result follows by substituting the bounds for κ and $d(p,q)$ into the expression for the width of the lune given by Lemma 9. This expression is an upper bound because increasing the radius (equivalently, decreasing κ) decreases the width of the lune. □

The proof of Lemma 9 gives an expression for Δ. The following lemma gives a lower bound on Δ in terms of the local feature size.

Lemma 16. *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface. Let $p, q \in \Sigma$ be points such that $d(p,q) \leq \rho\,\mathrm{lfs}(p)$ for some $\rho < 2\sqrt{5} - 4 \doteq 0.472136$. Let $h = \mathrm{Span}\{n_p, \vec{pq}\}$ be a plane. Parameterize the curve $h \cap \Sigma$ as $\gamma : [0,1] \to \Sigma$ such that $\gamma(0) = p$, $\gamma(1) = q$, and γ is a bijection from $[0,1]$ to $h \cap \Sigma$. Let B be the diametric ball of the segment pq, and suppose that $\gamma \subset B$. Let κ be the maximum curvature of γ. Let $C_{1/\kappa}$ and $C_{-1/\kappa}$ be the circles on h defined in Lemma 10, both of which enclose γ. Let Δ be the distance from the center of $C_{1/\kappa}$ (or the center of $C_{-1/\kappa}$) to the center of B. Then*

$$\Delta \geq \sqrt{1 - 2\rho - \frac{\rho^2}{4}}\,\mathrm{lfs}(p).$$

Proof. Substituting the upper bound on κ derived in Lemma 14 into the expression for Δ derived in Lemma 9 gives

$$\begin{aligned}
\Delta &= \frac{\sqrt{1 - \kappa^2\, d(p,q)^2/4}}{\kappa} \\
&\geq \sqrt{1-2\rho}\sqrt{1 - \frac{1}{\mathrm{lfs}(p)^2(1-2\rho)}\frac{\rho^2 \mathrm{lfs}(p)^2}{4}}\,\mathrm{lfs}(p) \\
&= \sqrt{1 - 2\rho - \frac{\rho^2}{4}}\,\mathrm{lfs}(p).
\end{aligned}$$

□

Combining Lemmas 11, 14, and 16 gives an upper bound on the distance between any point on γ and the nearest endpoint of the segment pq.

Lemma 17. *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface. Let $p, q \in \Sigma$ be points on Σ such that $d(p,q) \leq \rho\,\mathrm{lfs}(p)$ for some $\rho < 2\sqrt{5} - 4 \doteq 0.472136$. Let $h = \mathrm{Span}\{n_p, \vec{pq}\}$ be a plane. Parameterize the curve $h \cap \Sigma$ as $\gamma : [0,1] \to \Sigma$ such that $\gamma(0) = p$, $\gamma(1) = q$, and γ is a bijection from $[0,1]$ to $h \cap \Sigma$. Let B be the diametric ball of the segment pq, and suppose that $\gamma \subset B$. Let κ be the maximum curvature of γ. Let $C_{1/\kappa}$ and $C_{-1/\kappa}$ be the circles on h defined in Lemma 10, both of which enclose γ. Then for every point $x \in \gamma$,*

$$\min\{d(x,p), d(x,q)\} \leq \beta\,\mathrm{lfs}(p), \quad \text{where } \beta = \sqrt{2 - 4\rho - \sqrt{(1-2\rho)(4 - 8\rho - \rho^2)}}.$$

Proof. By Lemma 10, $\gamma \subset B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$, so $x \in B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$. The maximum distance between any point $y \in B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$ and the nearest of p or q is achieved where the bisector of pq intersects the boundary of $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$. By Lemma 11 (with $\theta = \pi/2$), this distance is $\sqrt{r^2 + \Delta^2 - 2r\Delta + d(p,q)^2/4}$, where $r = 1/\kappa$ and Δ is the distance from the center of $C_{1/\kappa}$ (or the center of $C_{-1/\kappa}$) to the center of B. By Lemma 14, $\kappa \leq 1/\left(\sqrt{1-2\rho}\,\mathrm{lfs}(p)\right)$. By Lemma 16, $\Delta \geq \sqrt{1 - 2\rho - \rho^2/4}\,\mathrm{lfs}(p)$. The result follows by substitution. □

3.4 Restricted Constrained Delaunay Triangulations

To define the restricted constrained Delaunay triangulation, we first define the *extended restricted Voronoi diagram* (or just *extended Voronoi diagram* for short) on the extended surface $\widetilde{\Sigma}$. As in Section 3.2, we define a distance function $\widehat{d}(p, q)$ that is the Euclidean distance in $\mathbb{R}^3$ if p and q are visible to each other along a straight path in $\widetilde{X}$, or ∞ if they cannot see each other. For any $v \in V$, the *extended restricted Voronoi cell* of v is

$$\operatorname{Vor}|_{\widetilde{\Sigma}} v = \{x \in \widetilde{\Sigma} : \widehat{d}(x, v) \le \widehat{d}(x, u),\ \forall u \in V\}.$$

An *extended restricted Voronoi face* is any nonempty set of points found by taking the intersection of one or more extended restricted Voronoi cells. The *extended Voronoi diagram* $\operatorname{Vor}|_{\widetilde{\Sigma}} V$ is the cell complex containing all the extended restricted Voronoi cells and faces.

We define the *restricted constrained Delaunay triangulation* (restricted CDT) $\operatorname{Del}|_{\widetilde{\Sigma}} V$ to be the simplicial complex dual to the extended Voronoi diagram. That is, a simplex is in $\operatorname{Del}|_{\widetilde{\Sigma}} V$ if its vertices' extended restricted Voronoi cells have a nonempty mutual intersection.

Although in principle $\operatorname{Del}|_{\widetilde{\Sigma}} V$ could contain a tetrahedron, we will assume that we can perturb $\widetilde{\Sigma}$ infinitesimally so that it does not pass through the tetrahedron's circumcenter; then $\operatorname{Del}|_{\widetilde{\Sigma}} V$ contains no tetrahedra. Similarly, just as a standard Delaunay triangulation in the plane can be ambiguous if four vertices lie on a common circle, if V has four or more cocircular vertices then $\operatorname{Del}|_{\widetilde{\Sigma}} V$ might have triangles with intersecting interiors that are trying to cover the same area, but an infinitesimal perturbation of V can fix this problem. We omit further details of these "degenerate" configurations, as standard remedies work well in this context. Where necessary, we assume every extended Voronoi vertex has degree three.

In this section, we discuss several combinatorial properties of extended Voronoi diagrams and restricted CDTs. A primary assumption is that no segment is too long. Specifically, for each segment $s \in S$ with endpoints $p, q \in \Sigma$, we assume that $d(p, q) \le \rho \operatorname{lfs}(p)$ for some sufficiently small value of ρ.

The following theorem shows that the sites whose extended restricted Voronoi cells lie in part on an extrusion Σ_s^+ must lie on the side of the cutting plane h_s opposite Σ_s^+ (excepting the endpoints of s, which lie on h_s). Thus the restricted Voronoi vertices on Σ_s^+ dualize to restricted Delaunay triangles that are also on the side of h_s opposite Σ_s^+. This theorem simplifies computing the restricted CDT, because an algorithm only needs to look at sites in one halfspace when computing the portion of $\operatorname{Vor}|_{\widetilde{\Sigma}} V$ that lies on Σ_s^+.

Theorem 18 (Cutting Plane Theorem)**.** *Let $s \in S$ be a segment with endpoints $p, q \in V$ such that $d(p, q) \le \rho \operatorname{lfs}(p)$ for $\rho \le 0.47$. Consider a point $x \in \Sigma_s^+$ and a site $r \in V \setminus \{p, q\}$ such that $x \in \operatorname{Vor}|_{\widetilde{\Sigma}} r$. Then r is strictly on the side of h_s opposite Σ_s^+. (The symmetric claim also holds for any $x \in \Sigma_s^-$.)*

Proof. Recall that s's cutting plane h_s induces a portal curve $\zeta_s \subset h_s$ and a unit vector b_s normal to h_s, with b_s determining the direction of the extrusions. For any point z, let $z_\perp \in \mathbb{R}$ denote the signed distance from z to h_s, being positive if z is on the same side of h_s as Σ_s^+ and negative if z is on the same side as Σ_s^-. As x lies on an extrusion Σ_s^+, we can write $x = \bar{x} + x_\perp b_s$ for some point $\bar{x} \in \zeta_s$ and

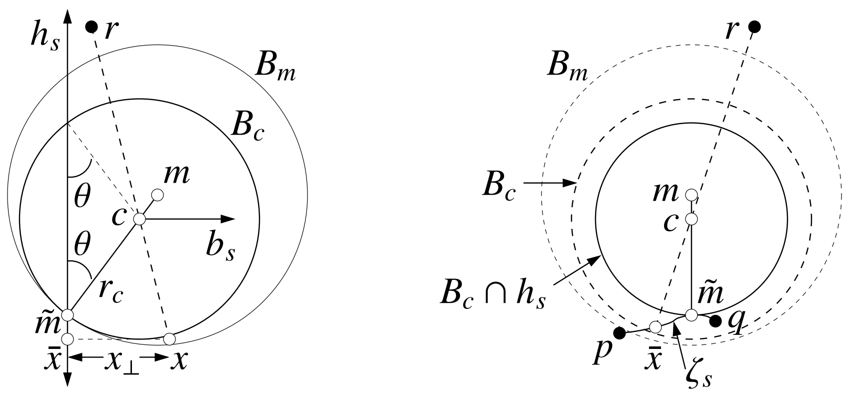

Figure 3.8: Left: side view of a cutting plane h_s with the extrusion Σ_s^+ extending to the right (with $\bar{x}x \subset \Sigma_s^+$). For a point $x \in \Sigma_s^+$ to lie in the Voronoi cell of a site r, the line segment rx must pass through the portal P_s at a point c, which lies on a line segment $m\tilde{m}$ connecting the center m of a medial ball B_m to a tangency point $\tilde{m} \in \Sigma$. (This figure is drawn so that c, m, and $\tilde{m}$ lie on the plane of the paper and h_s is orthogonal to the paper. However, r, x, and $\bar{x}$ do not generally lie on the plane of the paper.) As $r \notin B_m$ and $r_\perp \geq 0$ (that is, r is not left of h_s), the angle between the line $\overleftrightarrow{xcr}$ and the plane h_s is at most θ. Therefore, $x_\perp$ cannot be large. Right: another view of the same configuration; here h_s is the plane of the paper.

some scalar $x_\perp \geq 0$. (To apply this proof to $x \in \Sigma_s^-$, simply reverse the direction of b_s.) Observe that $\bar{x}$ is both the point on h_s closest to x and the point on ζ_s closest to x. By assumption, no site lies on ζ_s except p and q, so r cannot lie on ζ_s nor anywhere else on the portal P_s (as $r \in \Sigma$).

Suppose for the sake of contradicting Theorem 18 that $r_\perp \geq 0$ (i.e., r is on the same side of h_s as Σ_s^+ or $r \in h_s$). Our proof has two parts: first we show that $x_\perp \leq 0.17\,\mathrm{lfs}(p)$; then we show by different means that $x_\perp \geq 0.48\,\mathrm{lfs}(p)$. The theorem holds by contradiction. Both parts work by identifying a medial ball that touches ζ_s (a different ball for each part) that constrains the location of r, thereby constraining the location of x.

As $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} r$, there is a line segment $rx \subset \widetilde{X}$ that intersects the portal P_s at some point $c \in rx \cap P_s$ (because x lies on the extrusion Σ_s^+ and r is in the principal branch), illustrated in Figure 3.8. If there is more than one such intersection point, let c be the intersection point nearest r. Observe that $c_\perp \geq 0$ because $r_\perp \geq 0$ and $x_\perp \geq 0$. Recall that P_s is the union of the normal segments of the points on ζ_s. Let $\tilde{m} \in \zeta_s$ be the point whose normal segment $\ell_{\tilde{m}}$ contains c. Let m be the endpoint of $\ell_{\tilde{m}}$ such that $c \in \tilde{m}m$. Recall that m (like every normal segment endpoint) lies on Σ's medial axis M and is the center of an open medial ball B_m whose boundary is tangent to the surface Σ at $\tilde{m}$. As $B_m \cap \Sigma = \emptyset$, r cannot lie in B_m.

(Note that sometimes a medial "ball" is degenerate, effectively having infinite radius, thus being an open halfspace whose boundary is tangent to Σ at $\tilde{m}$. In this case, $\ell_{\tilde{m}}$ is a ray rather than a line segment, and although the point m is not defined, we can replace $\tilde{m}m$ with a ray originating at $\tilde{m}$ and lying on $\ell_{\tilde{m}}$. However, we note that a degenerate ball B_m cannot arise in this circumstance, because that would place the site r in the halfspace B_m. We won't prove that, because a degenerate B_m causes no difficulty for the rest of this proof.)

Let B_c be the open ball centered at c with $\tilde{m}$ on its boundary, and let $r_c = d(c, \tilde{m})$ be the radius of B_c, as illustrated in Figure 3.8. Observe that $B_c \subseteq B_m$, so B_c also does not intersect Σ nor contain r. Let θ be the angle at which the normal segment $\ell_{\tilde{m}}$ meets the cutting plane h_s, and observe that $c_\perp = r_c \sin\theta$. As $r \notin B_c$, $d(c,r) \geq r_c = c_\perp / \sin\theta$.

As $x \in \mathrm{Vor}|_{\tilde{\Sigma}} r$, $\widehat{d}(x,r) \leq \min\{d(x,p), d(x,q)\}$. As p, q, and $\bar{x}$ all lie on h_s, we have by Pythagoras' Theorem that $d(x,p)^2 = x_\perp^2 + d(\bar{x},p)^2$ and $d(x,q)^2 = x_\perp^2 + d(\bar{x},q)^2$. As $\bar{x} \in \zeta_s$, we have by Lemma 17 that $\min\{d(\bar{x},p), d(\bar{x},q)\} \leq \beta \,\mathrm{lfs}(p)$ where $\beta = \sqrt{2 - 4\rho - \sqrt{(1-2\rho)(4 - 8\rho - \rho^2)}}$. Hence

$$\begin{aligned} d(x,r)^2 &\leq \min\{d(x,p)^2, d(x,q)^2\} \\ &= x_\perp^2 + \min\{d(\bar{x},p)^2, d(\bar{x},q)^2\} \\ &\leq x_\perp^2 + \beta^2 \,\mathrm{lfs}(p)^2. \end{aligned} \tag{3.3}$$

We consider two cases: one for $r_\perp \leq c_\perp \leq x_\perp$ and one for $r_\perp \geq c_\perp \geq x_\perp$.

In the case where $r_\perp \leq c_\perp \leq x_\perp$ as illustrated in Figure 3.8, the angle between the vector b_s and the vector $\vec{cr}$ is at least 90°. As $r_\perp \geq 0$ and $r \notin B_c$, the angle between b_s and $\vec{cr}$ cannot exceed $90° + \theta$. Thus the angle between the line $\overleftrightarrow{xcr}$ and the plane h_s cannot exceed θ. So $x_\perp - c_\perp \leq d(x,c)\sin\theta$ and

$$\begin{aligned} \beta^2 \,\mathrm{lfs}(p)^2 &\geq d(x,r)^2 - x_\perp^2 \\ &= (d(x,c) + d(c,r))^2 - x_\perp^2 \\ &\geq \left(\frac{x_\perp - c_\perp}{\sin\theta} + \frac{c_\perp}{\sin\theta}\right)^2 - x_\perp^2 \\ &= \frac{x_\perp^2}{\sin^2\theta} - x_\perp^2 \\ &= x_\perp^2 \cot^2\theta. \end{aligned}$$

In the case where $r_\perp \geq c_\perp \geq x_\perp$,

$$\begin{aligned} \beta^2 \,\mathrm{lfs}(p)^2 &\geq d(x,r)^2 - x_\perp^2 \\ &\geq d(c,r)^2 - c_\perp^2 \\ &\geq \frac{c_\perp^2}{\sin^2\theta} - c_\perp^2 \\ &= c_\perp^2 \cot^2\theta \\ &\geq x_\perp^2 \cot^2\theta. \end{aligned}$$

In either case, $x_\perp \leq \beta \tan\theta \,\mathrm{lfs}(p)$. As $d(p,\tilde{m}) \leq d(p,q) \leq \rho \,\mathrm{lfs}(p)$, the Normal Variation Lemma (Lemma 21) implies that $\theta \leq \angle(n_p, n_{\tilde{m}}) \leq \eta(\rho)$ where $\eta(\delta) = \arccos\left(1 - \delta^2/(2\sqrt{1-\delta^2})\right)$. Hence $x_\perp \leq \beta \tan\eta(\rho) \,\mathrm{lfs}(p)$. (Note that β is a function of ρ.) From this inequality, it follows that $x_\perp \leq 0.17 \,\mathrm{lfs}(p)$ for all $\rho \in (0, 0.47]$, which completes the first part of our claim.

For the second part, we will show that $x_\perp \geq 0.48 \,\mathrm{lfs}(p)$, yielding a contradiction.

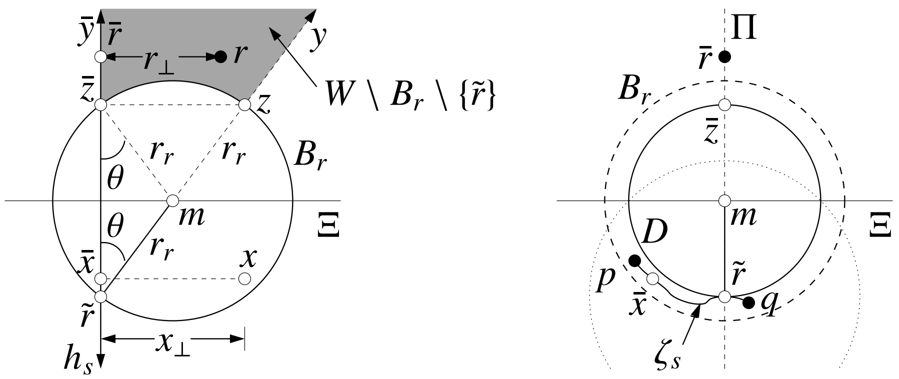

Figure 3.9: Left: side view of a cutting plane h_s. The paper is the plane Π orthogonal to the portal curve ζ_s at $\tilde{r}$, where $\tilde{r}$ is the point on ζ_s closest to the site r. Assuming that $r \in \Pi$, r must lie in the shaded region $W \setminus B_r \setminus y$. Note that x and $\bar{x}$ do not necessarily lie on the plane of the paper, but the other points do. For x to be in r's Voronoi cell rather than p's or q's, $x_\perp$ must be large. Right: another view of the same configuration; here the paper is the plane h_s, and we see side views of Π and Ξ.

Let $\tilde{r}$ be the point on ζ_s closest to r. As ζ_s is a smooth curve, there is a unique plane Π through $\tilde{r}$ that is locally orthogonal to ζ_s at $\tilde{r}$. As $\zeta_s \subset h_s$, Π is also orthogonal to the plane h_s. If $\tilde{r}$ is not an endpoint of ζ_s (p or q), then $r \in \Pi$, as $r\tilde{r}$ is orthogonal to ζ_s at $\tilde{r}$. However, if $\tilde{r}$ is an endpoint of ζ_s, then r might or might not lie on Π. If $r \notin \Pi$, then Π separates r from ζ_s.

Let $\ell_{\tilde{r}} \subset P_s$ be the normal segment through $\tilde{r}$, and let the line $L_{\tilde{r}}$ be the affine hull of $\ell_{\tilde{r}}$. Let $y \subset L_{\tilde{r}}$ be the ray that originates at $\tilde{r}$ and is half of $L_{\tilde{r}}$, on the positive side of h_s (the same side as Σ_s^+), as illustrated in Figure 3.9. As y is orthogonal to ζ_s at $\tilde{r}$, $y \subset \Pi$. Let the ray $\bar{y}$ be the orthogonal projection of y onto h_s. As the projection direction lies in Π, $\bar{y} \subset \Pi$. The rays y and $\bar{y}$ bound an infinite wedge $W \subset \Pi$ that has apex $\tilde{r}$, as illustrated.

As $x \in \mathrm{Vor}|_{\tilde{\Sigma}} r$, the site r is positioned so that the sightline rx enters the portal P_s from the correct side to access the extrusion Σ_s^+. Therefore, if $r \in \Pi$, then r lies in the wedge W because $r_\perp \geq 0$ (by assumption) but r lies on the same side of y as W. Note that $r \neq \tilde{r}$ and $r \notin y$ because $r \notin P_s$. On the other hand, if $r \notin \Pi$, let w be the point where the line segment rx intersects Π. Observe that $w_\perp \geq 0$ because $r_\perp \geq 0$ and $x_\perp \geq 0$. Hence, w lies in the wedge W.

The ray y passes through an endpoint m of the normal segment $\ell_{\tilde{r}}$, and $m \in M$ is the center of an open medial ball B_r that is tangent to Σ at $\tilde{r}$. As $B_r \cap \Sigma = \emptyset$, r cannot lie in B_r. (Note that the medial ball B_r cannot degenerate to a halfspace because then B_r would include $W \setminus \{\tilde{r}\}$, contradicting $r \in W \setminus B_r \setminus y$.)

Let $r_r = d(m, \tilde{r})$ be the radius of B_r. Observe that $\mathrm{lfs}(\tilde{r}) \leq r_r$. The rays y and $\bar{y}$ each intersect the boundary of B_r at two points: they both enter the ball at the ray origin $\tilde{r}$ and they each exit at another point. Let z be the point where y exits, and let $\bar{z}$ be the point where $\bar{y}$ exits, as illustrated in Figure 3.9. By circle geometry, $\angle z\bar{z}\tilde{r} = 90°$, so $\bar{z}$ is the point closest to z on h_s. Let θ be the wedge angle at which y meets $\bar{y}$. Then $d(\tilde{r}, z) = 2r_r$ and $d(\tilde{r}, \bar{z}) = 2r_r \cos\theta$.

As $\bar{x}$ and $\tilde{r}$ both lie on ζ_s, $d(\bar{x}, \tilde{r}) \leq d(p, q) \leq \rho \operatorname{lfs}(p)$ and $d(p, \tilde{r}) \leq d(p, q) \leq \rho \operatorname{lfs}(p)$. The Feature Translation Lemma (Lemma 38) implies that $\operatorname{lfs}(p) \leq \operatorname{lfs}(\tilde{r})/(1-\rho)$ and the Normal Variation Lemma (Lemma 21) implies that $\theta \leq \angle(n_p, n_{\tilde{r}}) \leq \eta(\rho)$.

Consider the open disk $D = B_r \cap h_s$, illustrated at right in Figure 3.9. The line segment $\tilde{r}\bar{z}$ is a diameter of D with length $2r_r \cos\theta$. D and $\bar{x}$ both lie on the cutting plane h_s but $\bar{x} \notin D$, as D does not intersect Σ. Hence we have by circle geometry that $\angle\tilde{r}\bar{x}\bar{z} \leq 90°$ and by Pythagoras' Theorem that $d(\bar{x}, \tilde{r})^2 + d(\bar{x}, \bar{z})^2 \geq d(\tilde{r}, \bar{z})^2 = 4r_r^2 \cos^2\theta$.

Let Ξ be the plane that bisects $\tilde{r}\bar{z}$; note that $m \in \Xi$ and Ξ cuts both B_r and D in half. We claim that $\tilde{r}$, x, and $\bar{x}$ all lie on the same side of Ξ; that is, $d(\tilde{r}, x) < d(\bar{z}, x)$ and $d(\tilde{r}, \bar{x}) < d(\bar{z}, \bar{x})$. This claim holds because for all $\rho \in [0, 0.53]$, $\rho < \sqrt{2}(1-\rho)\cos\eta(\rho)$ and therefore $d(\tilde{r}, \bar{x}) \leq d(p, q) \leq \rho \operatorname{lfs}(p) < \sqrt{2}(1-\rho)\operatorname{lfs}(p)\cos\eta(\rho) \leq \sqrt{2}\operatorname{lfs}(\tilde{r})\cos\theta \leq \sqrt{2}r_r\cos\theta = d(\tilde{r}, \bar{z})/\sqrt{2}$. This means that $\bar{x}$ is in the open ball with center $\tilde{r}$ and radius equal to $\sqrt{2}$ times the radius of D (see the dotted sphere at right in Figure 3.9). The boundary of this open ball intersects the boundary of D at two points that lie on Ξ, as illustrated. As $\bar{x} \notin D$, $\bar{x}$ must lie on the same side of Ξ as $\tilde{r}$. Furthermore, x must lie on the same side of Ξ as $\bar{x}$, because $\bar{x}$ is defined by projecting x in a direction parallel to Ξ.

We consider two cases, depending on whether $r \in \Pi$. For the case where $r \in \Pi$, r lies in the region $W \setminus B_r \setminus \{\tilde{r}\}$, which is shaded in Figure 3.9. The boundary of $W \setminus B_r \setminus \{\tilde{r}\}$ consists of three curves: a circular arc connecting z to $\bar{z}$, a ray that originates at z and is a subset of the ray y, and a ray $\bar{y}_{\bar{z}}$ that originates at $\bar{z}$ and is a subset of the ray $\bar{y}$. Let $\bar{r}$ be the point on h_s closest to r. As the figure makes clear, $\bar{r} \in \bar{y}_{\bar{z}}$. We see that $d(\bar{x}, \bar{r}) \geq d(\bar{x}, \bar{z})$, because $\bar{z}$ is the point in $\bar{y}_{\bar{z}}$ that is closest to Ξ (which is orthogonal to $\bar{y}_{\bar{z}}$) and $\bar{x}$ is on the other side of Ξ.

From Inequality (3.3) we have

$$\begin{aligned}
x_\perp^2 &\geq d(x, r)^2 - \beta^2 \operatorname{lfs}(p)^2 \\
&= d(\bar{x}, \bar{r})^2 + (x_\perp - r_\perp)^2 - \beta^2 \operatorname{lfs}(p)^2 \\
&\geq d(\bar{x}, \bar{z})^2 - \beta^2 \operatorname{lfs}(p)^2; \\
&\geq d(\tilde{r}, \bar{z})^2 - d(\bar{x}, \tilde{r})^2 - \beta^2 \operatorname{lfs}(p)^2 \\
&\geq 4r_r^2 \cos^2\theta - \rho^2 \operatorname{lfs}(p)^2 - \beta^2 \operatorname{lfs}(p)^2 \\
&\geq 4 \operatorname{lfs}(\tilde{r})^2 \cos^2\theta - (\rho^2 + \beta^2) \operatorname{lfs}(p)^2 \\
&\geq \left(4(1-\rho)^2 \cos^2\eta(\rho) - \rho^2 - \beta^2\right) \operatorname{lfs}(p)^2.
\end{aligned}$$

From this inequality, it follows that $x_\perp \geq 0.74 \operatorname{lfs}(p)$ for all $\rho \in (0, 0.47]$.

Now consider the case where $r \notin \Pi$. In this case, the point closest to r on the portal curve ζ_s is an endpoint of ζ_s; suppose without loss of generality that it is $\tilde{r} = q$. Recall that the plane Π is orthogonal to ζ_s at q. The site r and the curve ζ_s are on opposite sides of Π (otherwise, q could not be the point closest to r). Hence, r and x are on opposite sides of Π or $x \in \Pi$.

Let ℓ_q be q's normal segment, let B_q be the open medial ball tangent to Σ at q whose center m satisfies $m_\perp \geq 0$, and let $r_q = d(q, m)$ be the radius of B_q. (Note that as $\tilde{r} = q$, ℓ_q is the same normal segment we have already been calling $\ell_{\tilde{r}}$, and B_q is the same ball we have been calling B_r.) We define W, z, $\bar{z}$, y, and $\bar{y}$ as before; see Figure 3.10. As $B_q \cap \Sigma = \emptyset$, r cannot lie in B_q.

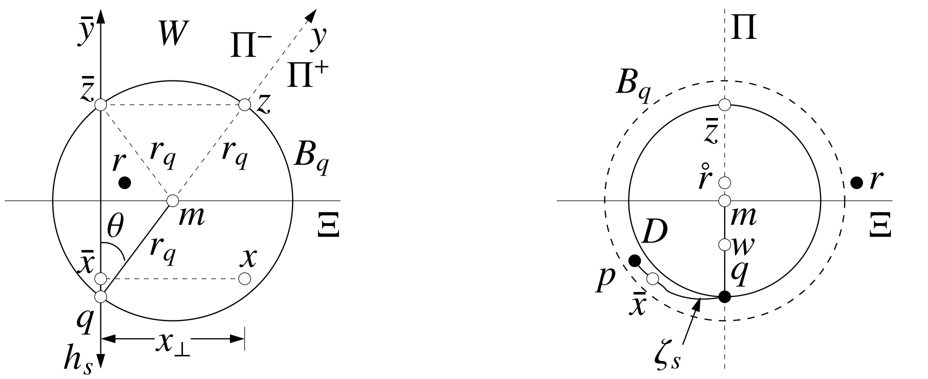

Figure 3.10: The circumstances of this figure are similar to those of Figure 3.9, but r does not lie on the plane Π (so it is not necessarily restricted to the region shaded in Figure 3.9).

Let the line $\ell \supset \ell_q$ be the affine hull of ℓ_q. The line ℓ cuts Π into two halfplanes; let Π^+ be the open halfplane on the positive side of ℓ, where the extrusion goes, and let Π^- be the closed halfplane on the negative side, so that $\bar{y} \subset \Pi^-$ and $W \subset \Pi^-$.

Define a coordinate system with q at the origin, such that for any point $p \in \mathbb{R}^3$, $p_\perp$ is p's signed distance from h_s (as we have already defined), p_Π is p's signed distance from Π, with the sign defined so that $r_\Pi > 0$ and $x_\Pi \leq 0$, and $p_\|$ is the coordinate in the direction $\bar{y}$ (the vertical axis in Figure 3.10). As $x \in \mathrm{Vor}\,|_{\widetilde{\Sigma}} r$, $d(x, r) \leq d(x, q)$. Therefore, $\|r\|^2 - 2x \cdot r \leq \|q\|^2 - 2x \cdot q$. We have chosen q as the origin, so we can write this as $\|r\|^2/2 \leq x \cdot r = x_\perp r_\perp + x_\| r_\| + x_\Pi r_\Pi$. As $x_\Pi r_\Pi \leq 0$, we can shorten this to $\|r\|^2/2 \leq x_\perp r_\perp + x_\| r_\|$. As $r \neq q$, this implies that $0 < x_\perp r_\perp + x_\| r_\|$.

As $r \notin B_q$, $d(m, r) \geq r_q$. Therefore, $\|m\|^2 - 2m \cdot r + \|r\|^2 \geq r_q^2$. But $\|m\| = r_q$, so we can write $\|r\|^2/2 \geq m \cdot r = m_\perp r_\perp + m_\| r_\|$. Combining this with an inequality from the previous paragraph gives $m_\perp r_\perp + m_\| r_\| \leq x_\perp r_\perp + x_\| r_\|$, or equivalently, $0 \leq (x_\perp - m_\perp) r_\perp + (x_\| - m_\|) r_\|$.

Let $\mathring{x}$ be the point closest to x on Π, and let $\mathring{r}$ be the point closest to r on Π. (Hence $\mathring{x}$ sets $\mathring{x}_\Pi = 0$ while retaining the coordinates $\mathring{x}_\perp = x_\perp$ and $\mathring{x}_\| = x_\|$.) We claim that $\mathring{x} \in \Pi^+$. (That is, with respect to the figure, $\mathring{x}$ is to the right of y.) To see this, suppose for the sake of contradiction that $\mathring{x} \in \Pi^-$. This implies that $\mathring{x}$ lies on or left of the ray $\vec{qm}$; equivalently, $m_\perp x_\| \geq m_\| x_\perp$. We know that $x_\perp \geq 0$, $m_\| > 0$, and $m_\perp > 0$, so it follows that $x_\| \geq 0$. Hence we can transform the inequality at the end of the last paragraph to $0 \leq (x_\perp x_\| - m_\perp x_\|) r_\perp + (x_\| - m_\|) x_\| r_\|$, and then to $0 \leq (x_\| - m_\|) x_\perp r_\perp + (x_\| - m_\|) x_\| r_\|$. As $m \in \Xi$ and x lies on the same side of Ξ as q, we have $x_\| - m_\| < 0$, so it follows that $0 \geq x_\perp r_\perp + x_\| r_\|$. But this contradicts the fact that $0 < x_\perp r_\perp + x_\| r_\|$ (from two paragraphs ago).

By this contradiction, we establish our claim that $\mathring{x} \in \Pi^+$. But the point $w = rx \cap \Pi$ must lie in Π^-; if it did not, the sightline from x to r would not exit the secondary space and enter the principal branch. Observe that w lies on the line segment $\mathring{r}\mathring{x}$, which implies that at least one of $\mathring{x}$ or $\mathring{r}$ is in Π^-. As $\mathring{x}$ is not, we must have $\mathring{r} \in \Pi^-$. Equivalently, $m_\perp r_\| \geq m_\| r_\perp$. As $r^+ \geq 0$, $\mathring{r}$ lies in the wedge W and thus $r_\| \geq 0$. Moreover, it is not possible that $\mathring{r} = q$, because then we would have $d(x, q) < d(x, r)$ and x could not be in r's Voronoi cell. Therefore, $\mathring{r}$ lies in $W \setminus \{q\}$ and thus $r_\| > 0$.

As $m_\| > 0$, we can transform the inequality from three paragraphs ago to $0 \leq (x_\perp - m_\perp) m_\| r_\perp + (x_\| -$

$m_\|)m_\| r_\|$, and then to $0 \leq (x_\perp - m_\perp)m_\perp r_\| + (x_\| - m_\|)m_\| r_\|$, and then to $0 \leq (x_\perp - m_\perp)m_\perp + (x_\| - m_\|)m_\| = x \cdot m - \|m\|^2$. (The last identity follows because $m_\Pi = 0$.)

In this coordinate system with origin q, $z = 2m$. We claim that $d(x, z) \leq d(x, q)$. To see this, observe that $d(x, z)^2 - d(x, q)^2 = \|z\|^2 - 2x \cdot z = 4\|m\|^2 - 4x \cdot m \leq 0$.

Therefore,

$$\begin{aligned}
0 &\geq d(x, z)^2 - d(x, q)^2 \\
&= d(\bar{x}, \bar{z})^2 + (x_\perp - z_\perp)^2 - d(\bar{x}, q)^2 - x_\perp^2 \\
&\geq d(\tilde{r}, \bar{z})^2 - d(\bar{x}, \tilde{r})^2 - 2x_\perp z_\perp + z_\perp^2 - \rho^2 \operatorname{lfs}(p)^2 \\
&\geq 4r_q^2 \cos^2\theta - \rho^2 \operatorname{lfs}(p)^2 - 4r_q x_\perp \sin\theta + 4r_q^2 \sin^2\theta - \rho^2 \operatorname{lfs}(p)^2 \\
&= 4r_q^2 - 2\rho^2 \operatorname{lfs}(p)^2 - 4r_q x_\perp \sin\theta \\
&\geq 4r_q^2 - 2\rho^2 \operatorname{lfs}(p)^2 - 4r_q x_\perp \sin\eta(\rho).
\end{aligned}$$

Therefore,

$$\begin{aligned}
x_\perp &\geq \frac{1}{\sin\eta(\rho)}\left(r_q - \frac{\rho^2}{2r_q}\operatorname{lfs}(p)^2\right) \\
&\geq \frac{1}{\sin\eta(\rho)}\left(\operatorname{lfs}(\tilde{r}) - \frac{\rho^2}{2\operatorname{lfs}(\tilde{r})}\operatorname{lfs}(p)^2\right) \\
&\geq \frac{1}{\sin\eta(\rho)}\left(1 - \rho - \frac{\rho^2}{2(1-\rho)}\right)\operatorname{lfs}(p).
\end{aligned}$$

From this inequality, it follows that $x_\perp \geq 0.48 \operatorname{lfs}(p)$ for all $\rho \in (0, 0.5]$. □

The next theorem shows that the restricted CDT $\mathrm{Del}|_{\overline{\Sigma}}V$ contains every edge in S.

Theorem 19. *Let $s \in S$ be a segment with endpoints $p, q \in V$ such that $d(p, q) \leq \rho \operatorname{lfs}(p)$ for $\rho \leq 0.47$. Then $\mathrm{Vor}|_{\overline{\Sigma}}p \cap \mathrm{Vor}|_{\overline{\Sigma}}q \neq \emptyset$. Hence pq is an edge in $\mathrm{Del}|_{\overline{\Sigma}}V$.*

Proof. We will show that $\mathrm{Vor}|_{\overline{\Sigma}}p$ meets $\mathrm{Vor}|_{\overline{\Sigma}}q$ on the extrusion Σ_s^+, as Figure 3.3 shows. (The same is true for Σ_s^-.) Let Π be the plane bisecting s. Consider the point $\bar{x} = \Pi \cap \zeta_s$ on the portal curve and the ray $\vec{x} = \Pi \cap \Sigma_s^+$, whose origin is $\bar{x}$. Let x be a point on $\vec{x}$, and note that $\bar{x}$ is the point closest to x on the portal plane h_s, and $x\bar{x}$ is perpendicular to $\bar{x}p$. We will show that for all $x \in \vec{x}$ sufficiently far from $\bar{x}$, $x \in \mathrm{Vor}|_{\overline{\Sigma}}p \cap \mathrm{Vor}|_{\overline{\Sigma}}q$.

Suppose that x lies in the Voronoi cell $\mathrm{Vor}|_{\overline{\Sigma}}r$ of a site $r \in V \setminus \{p, q\}$. Theorem 18 states that r is on the side of h_s opposite x. Therefore, there exists some $\delta > 0$ such that $d(x, r) \geq d(x, \bar{x}) + \delta$ for *every* site r (besides p and q) whose extended restricted Voronoi cell intersects $\vec{x}$. By Pythagoras' Theorem, $d(x, p)^2 = d(x, \bar{x})^2 + d(\bar{x}, p)^2$. So $x \in \mathrm{Vor}|_{\overline{\Sigma}}p$ if

$$\begin{aligned}
& d(x, p) \leq d(x, \bar{x}) + \delta \\
\leftrightarrow\ & d(x, \bar{x})^2 + d(\bar{x}, p)^2 \leq (d(x, \bar{x}) + \delta)^2 = d(x, \bar{x})^2 + 2\delta\, d(x, \bar{x}) + \delta^2 \\
\leftrightarrow\ & \frac{d(\bar{x}, p)^2 - \delta^2}{2\delta} \leq d(x, \bar{x}).
\end{aligned}$$

For any x that makes $d(x, \bar{x})$ sufficiently large, the last inequality holds, and hence $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} p$. As $d(x, p) = d(x, q)$, it also follows that $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} q$. Hence $\mathrm{Vor}|_{\widetilde{\Sigma}} p \cap \mathrm{Vor}|_{\widetilde{\Sigma}} q \neq \emptyset$. □

The shape of our extrusions Σ_s^+ and Σ_s^- is motivated in part by Theorem 19, which justifies the word "constrained" in "restricted constrained Delaunay triangulation."

The following theorem shows that the sites whose extended restricted Voronoi cells lie in part on an extrusion Σ_s^+ must lie in a ball centered on the midpoint of the segment s. The ball's radius is a little larger than the radius of s. This is a boon for efficiently computing the restricted CDT, because an algorithm needs to look only at sites near s when computing the portion of $\mathrm{Vor}|_{\widetilde{\Sigma}} V$ that lies on Σ_s^+.

Theorem 20 (Possession Theorem)**.** *Let $s \in S$ be a segment with endpoints $p, q \in V$ such that $d(p, q) \leq \rho \,\mathrm{lfs}(p)$ for $\rho \leq 0.47$. Let c be the midpoint of s. Let $x \in \Sigma_s^+$ (or $x \in \Sigma_s^-$). Let $r \in \widetilde{X}$ be a point that is not on the same side of the cutting plane h_s as x (though either point may lie on h_s) such that $d(x, r) \leq \min\{d(x, p), d(x, q)\}$. (Note that by Theorem 18, these conditions hold for any site $r \in V$ such that $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} r$). Then r lies in the ball $B(c, \lambda \,\mathrm{lfs}(p))$ with center c and radius $\lambda \,\mathrm{lfs}(p)$, where*

$$\lambda = \sqrt{1 - 2\rho}\left(1 - \sqrt{1 - \frac{\rho^2}{4\,(1 - 2\rho)}}\right) + \sqrt{(2 - 4\rho)\left(1 - \sqrt{1 - \frac{\rho^2}{4\,(1 - 2\rho)}}\right)}.$$

Proof. Let $\bar{x}$ be the point nearest to x on h_s, and observe that $\bar{x} \in \zeta_s$. Consider a closed ball B_x centered at x with radius $d(x, r)$. As $d(x, r) \leq \min\{d(x, p), d(x, q)\}$, neither p nor q is in the interior of B_x. The intersection of B_x with h_s is a closed disk with center $\bar{x}$. As $p, q \in h_s$, the radius of the disk $B_x \cap h_s$ is at most $\min\{d(\bar{x}, p), d(\bar{x}, q)\}$. As r and x are not on the same side of h_s, r lies in the ball centered at $\bar{x}$ with the same radius as $B_x \cap h_s$.

It follows that r must be included in the union of balls constructed by centering a ball at each point $w \in \zeta_s$ with radius $\min\{d(w, p), d(w, q)\}$. Lemma 10 states that ζ_s is a subset of the lune $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$, where κ is the curvature of ζ_s. (See Figure 3.7.) Let $\mathcal{B}$ be the union of all balls centered at every point in the lune. $\mathcal{B}$ is included in a ball centered at the midpoint c of pq with radius λ. To see this, we show that the point farthest from c in $\mathcal{B}$ is at most distance $\lambda \,\mathrm{lfs}(p)$ away. To find the farthest point, we focus on the balls centered at points on the boundary of the lune. By symmetry we need consider only one of the two arcs of the lune, and only the portion from p to the point directly above c in Figure 3.7.

We define a local coordinate system as shown in Figure 3.7, with the origin at the center of $C_{-1/\kappa}$, $c = (0, \Delta)$, and $p = (-d(p, q)/2, \Delta)$. In this coordinate system, we can parameterize $C_{-1/\kappa}$ as $C_{-1/\kappa}(\theta) = \sqrt{d(p, q)^2/4 + \Delta^2}(-\cos\theta, \sin\theta)$. The distance from c to the farthest point of each ball along the lune can be expressed as the sum of two distances, the distance from c to $C_{-1/\kappa}(\theta)$ plus the radius of the ball at $C_{-1/\kappa}(\theta)$.

$$d(c, C_{-1/\kappa}(\theta)) + d(C_{-1/\kappa}(\theta), p) = \sqrt{\frac{d(p,q)^2}{4} + 2\Delta^2 - 2\Delta\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\sin\theta}$$
$$+ \sqrt{\frac{d(p,q)^2}{2} + 2\Delta^2 - d(p,q)\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\cos\theta - 2\Delta\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\sin\theta}.$$

The first derivative of this sum with respect to θ is

$$-\frac{\Delta\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\cos\theta}{\sqrt{\frac{d(p,q)^2}{4} + 2\Delta^2 - 2\Delta\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\sin\theta}}$$
$$+\frac{\frac{1}{2}d(p,q)\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\sin\theta - \Delta\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\cos\theta}{\sqrt{\frac{d(p,q)^2}{2} + 2\Delta^2 - d(p,q)\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\cos\theta - 2\Delta\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\sin\theta}}.$$

By symmetry, we are interested in the zeros of the derivative in the range $\theta \in [\pi/2 - \arctan(d(p,q)/(2\Delta)), \pi/2]$. The derivative has a zero at $\theta = \pi/2 - \arctan(d(p,q)/(2\Delta))$, where $C_{-1/\kappa}(\theta) = p$, and no others in the range $[\pi/2 - \arctan(d(p,q)/(2\Delta)), \pi/2]$. Furthermore, the function is at a minimum at $\theta = \pi/2 - \arctan(d(p,q)/(2\Delta))$, with value $d(p,q)/2$. The derivative is positive at every other point in the range $[\pi/2 - \arctan(d(p,q)/(2\Delta)), \pi/2]$. Thus the maximum is achieved at one of the limits of the range, when $\theta = \pi/2$.

All that remains is to bound the sum of the distances when $\theta = \pi/2$. At $\theta = \pi/2$, Lemma 15 states that

$$d(c, C_{-1/\kappa}(\pi/2)) \leq \sqrt{1-2\rho}\left(1 - \sqrt{1 - \frac{\rho^2}{4(1-2\rho)}}\right)\operatorname{lfs}(p).$$

With the bound $d(c,p) = d(p,q)/2 \leq \rho\operatorname{lfs}(p)/2$, Pythagoras' Theorem implies that

$$\begin{aligned} d(C_{-1/\kappa}(\pi/2), p) &= \sqrt{d(c,p)^2 + d(c, C_{-1/\kappa}(\pi/2))^2} \\ &\leq \sqrt{\frac{\rho^2}{4} + (1-2\rho)\left(1 - \sqrt{1 - \frac{\rho^2}{4\,(1-2\rho)}}\right)^2}\operatorname{lfs}(p) \\ &= \sqrt{(2-4\rho)\left(1 - \sqrt{1 - \frac{\rho^2}{4\,(1-2\rho)}}\right)}\operatorname{lfs}(p). \end{aligned}$$

So $\mathcal{B}$ is a subset of a ball centered at c with radius $d(c, C_{-1/\kappa}(\pi/2)) + d(C_{-1/\kappa}(\pi/2), p) \leq \lambda\operatorname{lfs}(p)$. □

The Taylor series expansion around $\rho = 0$ gives $\lambda \approx \rho/2 + \rho^2/8 + 9\rho^3/64 + O(\rho^4)$. Thus the set of vertices whose extended Voronoi cells can intersect $\Sigma_s^\pm$ is included in a ball that is only somewhat larger than the diametric ball of s.

The next theorem shows that the 1-skeleton of the portion of the extended restricted Voronoi diagram on an extrusion Σ_s^+ is a tree. This implies several good things: the extrusion is subdivided into topological disks, and the combinatorial triangulation dual to that portion of the Voronoi diagram has the topology of a disk with s on its boundary. The theorem requires the following standard lemma with a recently improved bound.

Lemma 21 (Normal Variation Lemma [57])**.** *Let $\Sigma \subset \mathbb{R}^3$ be a bounded, smooth 2-manifold without boundary. Consider two points $p, q \in \Sigma$ and let $\delta = d(p, q)/\mathrm{lfs}(p)$. Let n_p and n_q be outward-directed vectors normal to Σ at p and q, respectively. If $\delta < \sqrt{4\sqrt{5} - 8} \doteq 0.9717$, then $\angle(n_p, n_q) \leq \eta(\delta)$ where*

$$\eta(\delta) = \arccos\left(1 - \frac{\delta^2}{2\sqrt{1-\delta^2}}\right) \approx \delta + \frac{7}{24}\delta^3 + O(\delta^5).$$

Theorem 22. *Consider a segment s with endpoints $p, q \in \Sigma$ such that $d(p, q) \leq \rho\,\mathrm{lfs}(p)$ for $\rho \leq 0.3368$ (implying $\lambda < 0.1957$, for λ defined as in Theorem 20). Let T be the 1-skeleton of $\mathrm{Vor}\,|_{\Sigma_s^+} V$ (or $\mathrm{Vor}\,|_{\Sigma_s^-} V$), the extended Voronoi diagram restricted to Σ_s^+ (or Σ_s^-). T is a tree.*

Proof. Suppose for the sake of contradiction that T contains a cycle. All the points inside and on the cycle lie on the extrusion Σ_s^+. Let v be a site whose extended Voronoi cell $\mathrm{Vor}\,|_{\tilde{\Sigma}} v$ has points inside the cycle. Let L_v be the line normal to Σ at v. Among all the points of $\mathrm{Vor}\,|_{\tilde{\Sigma}} v$ that are on or inside the cycle, let $x \in \Sigma_s^+$ be the one that is closest to L_v (breaking ties arbitrarily). Let h_s be the cutting plane for s, and let $\zeta_s \subset h_s \cap \Sigma$ be the portal curve for s. Let $\bar{x} \in \zeta_s$ be the orthogonal projection of x onto h_s. Let n_x be a vector normal to Σ_s^+ at x. Let T_x be the plane tangent to Σ_s^+ at x; thus n_x is normal to T_x. The ray $\vec{x}$ originating at $\bar{x}$ and passing through x is a subset of the extrusion Σ_s^+ and a subset of T_x.

Let $W \subseteq V$ be the set of sites (including v) whose extended Voronoi cells contain x. Every site in W is equidistant from x. Consider the standard, three-dimensional Voronoi diagram of W and let P_v be v's Voronoi cell in that diagram; P_v is an unbounded convex polyhedron. If we treat x as if it were in $\mathbb{R}^3$ (rather than in a secondary branch), then x lies in every Voronoi cell of that diagram, including P_v.

Let $B, B' \subset \mathbb{R}^3$ be the two open balls of radius $\mathrm{lfs}(v)$ tangent to Σ at v, and let $o, o' \in \mathbb{R}^3$ be their centers, respectively. Neither ball intersects Σ, so neither ball contains any site in W. Hence $d(v, o) \leq d(w, o)$ and $d(v, o') \leq d(w, o')$ for all $w \in W$, hence $o \in P_v$ and $o' \in P_v$. As P_v is convex, $\triangle oo'x \subset P_v$ (again treating x as if it were in $\mathbb{R}^3$). As v is strictly in the interior of P_v and $v \in oo'$, the relative interior of $\triangle oo'x$ is a subset of the interior of P_v.

Let c be the center of the segment s. By Theorem 20, $d(v, c) \leq \lambda\,\mathrm{lfs}(p)$, so $d(v, p) \leq d(v, c) + d(c, p) \leq (\lambda + \rho/2)\,\mathrm{lfs}(p)$. Likewise, $d(v, \bar{x}) \leq (\lambda + \rho/2)\,\mathrm{lfs}(p)$. By the Feature Translation Lemma (Lemma 38), $d(v, \bar{x}) \leq (\lambda + \rho/2)\,\mathrm{lfs}(v)/(1 - \lambda - \rho/2)$. Also by the Feature Translation Lemma and the fact that $d(p, \bar{x}) \leq \rho\,\mathrm{lfs}(p)$, $d(v, \bar{x}) \leq (\lambda + \rho/2)\,\mathrm{lfs}(\bar{x})/(1 - \rho)$. By the Normal Variation Lemma (Lemma 21), $\angle(n_v, n_p) \leq \eta(\lambda + \rho/2)$ and $\angle(n_v, n_{\bar{x}}) \leq \eta((\lambda + \rho/2)/(1 - \rho))$.

To establish an upper bound on $\angle(n_v, n_x)$, we divide the variation into portions transverse to h_s and orthogonal to h_s. The former portion is equal to the transverse portion of $\angle(n_v, n_{\bar{x}})$. The latter portion is equal to the orthogonal portion of $\angle(n_v, n_p)$. Therefore, $\angle(n_v, n_x) \leq \angle(n_v, n_{\bar{x}}) + \angle(n_v, n_p) \leq \eta((\lambda + \rho/2)/(1 - \rho)) + \eta(\lambda + \rho/2)$.

The ball centers o and o' lie on opposite sides of T_x, by the following reasoning. The upper bound on $d(v, \bar{x})$ given above implies an upper bound on the angles $\angle vo\bar{x}$ and $\angle vo'\bar{x}$. Specifically, let $\theta = \angle vo\bar{x}$; as $v, \bar{x} \notin B$, we have $d(v, \bar{x}) \geq 2\,\mathrm{lfs}(v)\sin(\theta/2)$, so $\theta \leq 2\arcsin(d(v, \bar{x})/(2\,\mathrm{lfs}(v)))$. Symmetrically, as $v, \bar{x} \notin B'$, the same inequality holds with θ replaced by $\theta' = \angle vo'\bar{x}$. Recall that T_x passes through $\bar{x}$ with normal vector n_x, and that n_v is parallel to vo and vo'. Therefore, if $\angle(n_v, n_x) < 90° - 2\arcsin(d(v, \bar{x})/(2\,\mathrm{lfs}(v)))$, then o and o' lie on opposite sides of T_x. This inequality must hold if $\eta((\lambda + \rho/2)/(1 - \rho)) + \eta(\lambda + \rho/2) < 90° - 2\arcsin((\lambda + \rho/2)/(2 - 2\lambda - \rho))$, and the latter inequality holds for all $\rho \in [0, 0.3368]$, so o and o' lie on opposite sides of T_x.

By Theorem 18, the site v is on the side of h_s opposite from Σ_s^+. We claim that $x \notin L_v$. To see that, observe that $d(x, \bar{x}) < d(x, v)$, as $\bar{x}$ is the point closest to x on h_s whereas v is on the opposite side of h_s. Recall that $\bar{x} \in \Sigma$. Hence x does not lie on v's normal segment ℓ_v (if it did, that would imply that $d(x, v) \leq d(x, \bar{x})$). Observe that $o, o' \in \ell_v \subset L_v$. As o and o' lie on opposite sides of T_x, x cannot lie on any other part of L_v either. Hence $x \notin L_v$ and $\triangle oo'x$ is a nondegenerate triangle.

By Lemma 46, the line segment vx intersects no portal boundary; therefore, there is an open neighborhood N of x in Σ_s^+ sufficiently small that every point in N is visible from v. Moreover, Lemma 46 implies that every extended Voronoi cell restricted to Σ_s^+ is closed; hence, there is an open neighborhood of x in Σ_s^+ that intersects no extended Voronoi cell of a site not in W. Let $N \subset \Sigma_s^+$ be an open neighborhood of x that satisfies both properties. Every point in $P_v \cap N$ is in $\mathrm{Vor}|_{\bar{\Sigma}} v$, and every such point in the interior of P_v is in the relative interior of $\mathrm{Vor}|_{\bar{\Sigma}} v$. Recall that $\triangle oo'x \subset P_v$ and the relative interior of $\triangle oo'x$ is a subset of the interior of P_v. As N is smooth and the line segments ox and $o'x$ are on opposite sides of the plane T_x tangent to N at x, $\triangle oo'x \cap N$ includes a path that begins at x and immediately enters the relative interior of $\triangle oo'x$, and therefore enters the interior of P_v, and therefore enters the relative interior of $\mathrm{Vor}|_{\bar{\Sigma}} v$. As $o, o' \in L_v$, every point on the path except x is closer to L_v than x is. This contradicts the assumption that x is the point of $\mathrm{Vor}|_{\bar{\Sigma}} v$ on or inside the cycle that is closest to L_v. Hence, by contradiction, T does not contain a cycle. □

3.5 Topological Guarantees

Here we introduce some conditions in which a restricted CDT is homeomorphic to the original surface Σ, with a view to applications in guaranteed-quality surface mesh generation. The *nearest point map* ν maps a point $x \in \mathbb{R}^d \setminus M$ to the point $\nu(x)$ nearest x on Σ. We show that the nearest point map, with its domain restricted to the underlying space of the restricted CDT $\mathrm{Del}|_{\bar{\Sigma}} V$, is a homeomorphism from the underlying space to the surface Σ.

Our proof that ν is a homeomorphism has three conditions: a segment length condition, that each segment $s \in S$ with endpoints p and q satisfies $d(p, q) \leq 0.3368\,\mathrm{lfs}(p)$; a *sampling condition* requiring the sites V to be sufficiently dense; and an *encroachment condition* that prevents vertices in V from being too close to a segment, to prevent the possibility of very skinny triangles. All three conditions could be enforced without difficulty by a mesh generation algorithm that inserts new vertices on Σ (sometimes subdividing segments).

To understand the sampling condition, consider a surface Σ without boundary, a finite vertex set $V \subset \Sigma$, a set of segments S whose endpoints are in V, and a set Z of portal curves containing one curve ζ_s for each $s \in S$. Recall the slitted surface $\overline{\Sigma}_S$, defined in Section 3.2 to be the completion of $\Sigma - \bigcup_{s \in S} \mathrm{Int}(\zeta_s)$. We say that V is a *constrained ϵ-sample of* (Σ, S, Z) if V contains every endpoint of every segment in S and for every point $x \in \overline{\Sigma}_S$, there is a site $v \in V$ such that $\widehat{d}(x, v) \leq \epsilon\, \mathrm{lfs}(x)$. That is, the ball centered at x with radius $\epsilon\, \mathrm{lfs}(x)$ contains at least one sample point visible from x. Here, visibility and $\widehat{d}$ are as defined in Section 3.2, and they are what differentiates a constrained ϵ-sample from a standard ϵ-sample. (If S is empty, the two are the same.) Our homeomorphism proof requires that V be a constrained 0.3202-sample of (Σ, S, Z).

The encroachment condition applies only to restricted Delaunay triangles whose dual extended Voronoi vertices lie on an extrusion. Let τ be such a triangle. The *circumradius* r of τ is the radius of the unique circle that passes through τ's three vertices. Let v be the vertex of τ at τ's largest plane angle. We require that $r \leq 0.3606\, \mathrm{lfs}(v)$. The purpose of this restriction is to prevent the existence of "inverted" triangles in $\mathrm{Del}\,|_{\overline{\Sigma}} V$, which create "foldovers" in the map from $\mathrm{Del}\,|_{\overline{\Sigma}} V$ to Σ. With foldovers, the nearest point map is not injective and therefore not a homeomorphism.

The sampling and encroachment conditions both rule out triangles with circumradii that are excessively large relative to the local feature size. A large circumradius implies either that the triangle is large, or that it has a large plane angle (close to 180°). Imposing these conditions is consistent with a mesh generator's goal of producing only well-shaped triangles, so the conditions are not onerous. Nevertheless, there are other applications such as surface reconstruction where the encroachment condition is not a natural condition. The restricted CDT may nevertheless still be useful in that context; see the Conclusions for speculations.

Our main result is the following theorem. Unfortunately, the proof is fifteen pages long, so we delay it to Appendix A.3. To keep the proof from being even longer, we impose a nondegeneracy condition that every extended Voronoi vertex has degree three, which can be enforced by infinitesimal perturbations of Σ and V (but isn't necessary in practice).

Theorem 23. *Let V be a constrained ϵ-sample of (Σ, S, Z) for some $\epsilon \leq 0.3202$. Suppose that for every segment $pq \in S$, $d(p, q) \leq 0.3368\, \mathrm{lfs}(p)$. Suppose that every extended Voronoi vertex in $\mathrm{Vor}\,|_{\overline{\Sigma}} V$ has degree three. Suppose that for every extended Voronoi vertex that lies on an extrusion, its dual restricted Delaunay triangle satisfies $r \leq 0.3606\, \mathrm{lfs}(v)$, where r is τ's circumradius and v is the vertex of τ at τ's largest plane angle. Then the nearest point map $\nu : |\mathrm{Del}\,|_{\overline{\Sigma}} V| \to \Sigma$ is a homeomorphism.*

We sketch the main ideas of the proof. We call $\overline{\Sigma}_S$ the *principal surface*; recall from Section 3.2 that $\overline{\Sigma}_S$ is the topological space we obtain by cutting slits in Σ and completing the space, but before gluing on the extrusions. Each site $v \in V$ has a *principal Voronoi cell* $\mathrm{Vor}\,|_{\overline{\Sigma}_S} v = \overline{\Sigma}_S \cap \mathrm{Vor}\,|_{\overline{\Sigma}} v$, which excludes the portion of the extended restricted Voronoi cell on the extrusions. We call an extended restricted Voronoi vertex a *principal vertex* if it lies on $\overline{\Sigma}_S$, or a *secondary vertex* otherwise (i.e., it lies on an extrusion but not on a portal curve).

In Appendix A.3 we show that, as V is a constrained 0.3202-sample, each principal vertex dualizes to a triangle whose circumradius is not large (relative to the local feature size). The encroachment condition implies that each secondary vertex dualizes to a triangle whose circumradius is not large. The bounds on circumradii allow us to prove that the nearest point map restricted to any single restricted Delaunay triangle is a homeomorphism. Moreover, there is a sense in which the map preserves orientation: for any extended Voronoi vertex u whose dual extended Delaunay triangle is $\tau = \triangle pp'p''$ the sites p, p', and p'' are in counterclockwise order around $\nu(\tau)$ if and only if the cells $\mathrm{Vor}\,|_{\widetilde{\Sigma}} p$, $\mathrm{Vor}\,|_{\widetilde{\Sigma}} p'$, and $\mathrm{Vor}\,|_{\widetilde{\Sigma}} p''$ adjoin u in counterclockwise order around u (as seen from outside Σ). From this, we argue that along each of its edges, each restricted Delaunay triangle adjoins another restricted Delaunay triangle with a consistent orientation, and therefore the restricted Delaunay triangles must cover the whole surface Σ—that is, the nearest point map is a surjection from $|\,\mathrm{Del}\,|_{\widetilde{\Sigma}} V|$ to Σ.

Given a constrained 0.44-sample V, for any site $v \in V$ that does not adjoin a segment, its principal Voronoi cell $\mathrm{Vor}\,|_{\widetilde{\Sigma}_S} v$ is homeomorphic to a closed disk. Theorem 22 implies that each portion of $\mathrm{Vor}\,|_{\widetilde{\Sigma}} v$ on an extrusion is also homeomorphic to a closed disk (if we add a "point at infinity" to each extrusion to make it closed). It follows that the complete extended Voronoi cell $\mathrm{Vor}\,|_{\widetilde{\Sigma}} v$ is a topological disk, because it is a disk with other disks glued along portions of its boundary. Because the boundary of $\mathrm{Vor}\,|_{\widetilde{\Sigma}} v$ is a simple loop, the restricted Delaunay triangles adjoining v form a fan of triangles around v whose union is a topological disk. The argument is just a little more complicated for a site w that adjoins one or more segments: those segments may cut w's principal Voronoi cell into several disks, but still the restricted Delaunay triangles adjoining w form a fan around w. From that we prove that the nearest point map is an injection from $|\,\mathrm{Del}\,|_{\widetilde{\Sigma}} V|$ to Σ (there are no "foldovers" that cause any part of Σ to be covered by multiple triangles), and therefore a homeomorphism.

Chapter 4

On the Geometry of Adversarial Examples

4.1 Introduction

Deep learning at scale has led to breakthroughs on important problems in computer vision [60], natural language processing [96], and robotics [63]. Shortly thereafter, the interesting phenomena of *adversarial examples* was observed. A seemingly ubiquitous property of machine learning models where perturbations of the input that are imperceptible to humans reliably lead to confident incorrect classifications [87, 41]. What has ensued is a standard story from the security literature: a game of cat and mouse where defenses are proposed only to be quickly defeated by stronger attacks [7]. This has led researchers to develop methods which are provably robust under specific attack models [66, 94, 85, 72]. As machine learning proliferates into society, including security-critical settings like health care [38] or autonomous vehicles [24], it is crucial to develop methods that allow us to understand the vulnerability of our models and design appropriate counter-measures.

This chapter discusses a geometric framework for analyzing the phenomenon of adversarial examples. We leverage the observation that datasets encountered in practice exhibit low-dimensional structure despite being embedded in very high-dimensional input spaces. This property is colloquially referred to as the "Manifold Hypothesis": the idea that low-dimensional structure of 'real' data leads to tractable learning. We model data as being sampled from class-specific low-dimensional manifolds embedded in a high-dimensional space. We consider a threat model where an adversary may choose *any* point on the data manifold to perturb by ϵ in order to fool a classifier. In order to be robust to such an adversary, a classifier must be correct everywhere in an ϵ-tube around the data manifold. Observe that, even though the data manifold is a low-dimensional object, this tube has the same dimension as the entire space the manifold is embedded in. Our analysis argues that adversarial examples are a natural consequence of learning a decision boundary that classifies all points on a low-dimensional data manifold correctly, but classifies many points near the manifold incorrectly. The high *codimension*, the difference between the dimension of the data manifold and the dimension of the embedding space, is a key source of the pervasiveness of adversarial examples.

This chapter makes the following contributions.

- A geometric framework, inspired by the manifold reconstruction literature, that formalizes the manifold hypothesis described above and our attack model.

- We highlight the role *codimension* plays in vulnerability to adversarial examples. As the codimension increases, there are an increasing number of directions off the data manifold in which to construct adversarial perturbations. Prior work has attributed vulnerability to adversarial examples to input dimension [40].

- We apply this framework to prove the following results: (1) we show that the choice of norm to restrict an adversary is important in that there exists a tradeoff between being robust to different norms: we present a classification problem where improving robustness under the L_∞ norm requires a loss of $\Omega(1 - 1/\sqrt{d})$ in robustness to the L_2 norm; (2) we show that a common approach, training against adversarial examples drawn from balls around the training set, is insufficient to learn robust decision boundaries with realistic amounts of data; and (3) we show that nearest neighbor classifiers do not suffer from this insufficiency, due to geometric properties of their decision boundary away from data, and thus represent a potentially robust classification algorithm.

- A modification to the standard paradigm of adversarial training. We replace the L_p-ball constraint with the Voronoi cells of the training data, which have several advantages detailed in Section 4.8. In particular, we need not set the maximum perturbation size ϵ as part of the training procedure. The Voronoi cells adapt to the maximum allowable perturbation size locally on the data distribution. We show how to construct adversarial examples within the Voronoi cells and how to incorporate Voronoi constraints into standard adversarial training. In Section 4.10 we show that adversarial training with Voronoi constraints gives state-of-the-art robustness results on MNIST and competitive results on CIFAR10.

4.2 Related Work

This chapter approaches the problem of adversarial examples using techniques and intuition from the manifold reconstruction literature. Both fields have a great deal of prior work, so we focus on only the most related papers here.

Adversarial Examples

There has been a long line of work on the theory of adversarial examples. Schmidt et al. [76] explore the sample complexity required to produce robust models. They demonstrate a simple setting, a mixture of two Gaussians, in which a linear classifier with near perfect natural accuracy can be learned from a single sample, but *any* algorithm that produces *any* binary classifier requires $\Omega(\sqrt{d})$ samples to produce a robust classifier. Followup work by Bubeck et al. [16] suggests that adversarial examples may arise from computational constraints. They exhibit pairs of distributions that differ only in a k-dimensional subspace, and are otherwise standard Gaussians, and show that while it is information-theoretically possible to distinguish these distributions, it requires exponentially many queries in the statistical query model of computation. We note that both of these constructions produce distributions whose support is the entirety of $\mathbb{R}^d$. The work of Gilmer et al. [40] experimentally evaluated the setting of two concentric under-sampled 499-spheres embedded in $\mathbb{R}^{500}$, and concluded that adversarial examples occur on the data manifold. In contrast, we present a geometric framework for proving robustness guarantees for learning algorithms, that makes no assumptions on the decision boundary. We carefully sample the data manifold in order to highlight the importance of *codimension*; adversarial examples exist *even* when the manifold is perfectly classified. Additionally we explore the importance of the spacing between the constituent data manifolds, sampling requirements for learning algorithms, and the relationship between model complexity and robustness. Shafahi et al. [80] suggest that adversarial examples may be an unavoidable consequence of the high-dimensional geometry of data. Their result depends upon the use of an isopermetric inequality. The main drawback of these prior works is that they assume that the support of the data distribution has full or nearly full dimension. We do not believe this to be the case in practice, instead we believe that the data distribution is often supported on a very low-dimensional subset of $\mathbb{R}^d$.

Wang et al. [91] explore the robustness of k-nearest neighbor classifiers to adversarial examples. In the setting where the Bayes optimal classifier is uncertain about the true label of each point, they show that k-nearest neighbors is not robust if k is a small constant. They also show that if $k \in \Omega(\sqrt{dn \log n})$, then k-nearest neighbors is robust. Using our geometric framework we show a complementary result: in the setting where each point is certain of its label, 1-nearest neighbors is robust to adversarial examples.

The decision and medial axes defined in Section 4.3 are maximum margin decision boundaries. Hard margin SVMs define define a linear separator with maximum margin, maximum distance from the training data [26]. Kernel methods allow for maximum margin decision boundaries that are non-linear by using additional features to project the data into a higher-dimensional feature space [81]. The decision and medial axes generalize the notion of maximum margin to account for the arbitrary curvature of the data manifolds. There have been attempts to incorporate maximum margins into deep learning [86, 65, 64, 37], often by designing loss functions that encourage large margins at either the output [86] or at any layer [37]. In contrast, the decision axis is defined on the input space and we use it as an analysis tool for proving robustness guarantees.

Adversarial training, the process of training on adversarial examples generated in L_p-balls around the training data, is a very natural approach to constructing robust models and was originally proposed by Goodfellow et al. [41]. Madry et al. [66] formalized the adversarial training objective and highlighted the importance of a strong adversary for constructing adversarial examples in the inner training loop. Their approach to adversarial training, which utilized a projected gradient descent adversary, produced some of the first empirically robust models which were not later broken by stronger attacks. There's was the *only* approach surveyed by Athalye et al. [7] which was not either fully circumvented by [7] or in a later paper [51]. More recently, the celebrated algorithm TRADES [97] has been proposed, which attempts to provide a principled way to trade off between robustness and natural accuracy. The analysis that inspires TRADES decomposes the robust error into two terms: natural error and error near the decision boundary. The yields an objective function with two terms, one which encourages accuracy and another which pushes the decision boundary away from the data distribution. Constructing a decision boundary that is far from the data distribution is explored in other heuristic works such as [33, 50, 47]. The approach we describe in Section 4.9 falls into this class of defenses and so we will compare exclusively against such defenses.

The frequency with which heuristic defenses have been defeated by stronger attacks has led to a line of work on certifiable robustness, which can guarantee that there exists no perturbation within an L_p-ball of radius ϵ which causes the classifier to change its classification. One of the first works by Wong et al. [94] proposed to approximate the set of possible activations of every L_∞-bounded perturbation by propagating upper and lower bounds for each activation through the network. These upper and lower bounds are used to construct a convex outer approximation to the set of possible activations in the final layer, and a linear program is used to certify that this convex approximation does not intersect the decision boundary. This initial work had several notable drawbacks, and several subsequent works have attempted to improve upon these initial results [92, 67, 39, 95, 84]. However the fundamental problems have remained: (1) these approaches do not scale to larger networks despite considerable effort, (2) they often depend crucially on the specific details of the architecture, and (3) the size of ϵ which can be certified is often considerably smaller than what we observe to be empirically robust. A different approach to certified robustness which addresses some of these concerns, called randomized smoothing [61, 25], has recently been proposed. Randomized smoothing leverages the ability of *any* classifier f to perform well on Gaussian noise to construct a new classifier g which is certifiably robust under adversarial L_2 perturbations. Unlike prior approaches to certified robustness, randomized smoothing is a simple approach which does not depend on the architecture details of the classifier. Its main drawback is that it is currently limited to L_2. We also note that more recent work has combined randomized smoothing with adversarial training to produce even more certifiably robustness classifiers in L_2 [75]. Since the goal and limitations of these method are often different from heuristic approaches we do not compare our method against these approaches.

Manifold Reconstruction

Manifold reconstruction is the problem of discovering the structure of a k-dimensional manifold embedded in $\mathbb{R}^d$, given *only* a set of points sampled from the manifold. A large vein of research in manifold reconstruction develops algorithms that are *provably good*: if the points sampled from the underlying manifold are sufficiently dense, these algorithms are guaranteed to produce a geometrically accurate representation of the unknown manifold with the correct topology. The output of these algorithms is often a *simplicial complex*, a set of simplices such as triangles, tetrahedra, and higher-dimensional variants, that approximate the unknown manifold. In particular these algorithms output subsets of the Delaunay triangulation, which along with their geometric dual the Voronoi diagram, have properties that aid in proving geometric and topological guarantees [36].

The field first focused on curve reconstruction in $\mathbb{R}^2$ [3] and subsequently in $\mathbb{R}^3$ [30]. Soon after algorithms were developed for surface reconstruction in $\mathbb{R}^3$, both in the noise-free setting [2, 6] and in the presence of noise [29]. We borrow heavily from the analysis tools of these early works, including the medial axis and the reach. However we emphasize that we have adapted these tools to the learning setting. To the best of our knowledge, our work is the first to consider the medial axis under different norms.

In higher-dimensional embedding spaces (large d), manifold reconstruction algorithms face the *curse of dimensionality*. In particular, the Delaunay triangulation, which forms the bedrock of algorithms in low-dimensions, of n vertices in $\mathbb{R}^d$ can have up to $\Theta(n^{\lceil d/2 \rceil})$ simplices. To circumvent the curse of dimensionality, algorithms were proposed that compute subsets of the Delaunay triangulation restricted to the k-dimensional tangent spaces of the manifold at each sample point [10]. Unfortunately, progress on higher-dimensional manifolds has been limited due to the presence of so-called "sliver" simplices, poorly shaped simplices that cause in-consistences between the local triangulations constructed in each tangent space [19, 10]. Techniques that provably remove sliver simplices have prohibitive sampling requirements [22, 10]. Even in the special case of surfaces ($k = 2$) embedded in high dimensions ($d > 3$), algorithms with practical sampling requirements have only recently been proposed [58]. Our use of tubular neighborhoods as a tool for analysis is borrowed from [31] and [58].

In this chapter we are interested in *learning* robust decision boundaries, *not* reconstructing the underlying data manifolds, and so we avoid the use of Delaunay triangulations and their difficulties entirely. In Section 4.5 we present robustness guarantees for two learning algorithms in terms of a sampling condition on the underlying manifold. These sampling requirements scale with the dimension of the underlying manifold k, *not* with the dimension of the embedding space d.

4.3 The Geometry of Data

We model data as being sampled from a set of low-dimensional manifolds (with or without boundary) embedded in a high-dimensional space $\mathbb{R}^d$. We use k to denote the dimension of a manifold $\mathcal{M} \subset \mathbb{R}^d$. The special case of a 1-manifold is called a *curve*, and a 2-manifold is a *surface*. The *codimension* of $\mathcal{M}$ is $d - k$, the difference between the dimension of the manifold and the dimension of the embedding space. The "Manifold Hypothesis" is the observation that in practice, data is often sampled from manifolds, usually of high codimension.

In this chapter we are primarily interested in the classification problem. Thus we model data as being sampled from C *class manifolds* $\mathcal{M}_1, \ldots, \mathcal{M}_C$, one for each class. When we wish to refer to the entire space from which a dataset is sampled, we refer to the *data manifold* $\mathcal{M} = \cup_{1 \leq j \leq C} \mathcal{M}_j$. We often work with a finite sample of n points, $X \subset \mathcal{M}$, and we write $X = \{X_1, X_2, \ldots, X_n\}$. Each sample point X_i has an accompanying class label $y_i \in \{1, 2, \ldots, C\}$ indicating which manifold $\mathcal{M}_{y_i}$ the point X_i is sampled from.

Consider a L_p-ball B centered at some point $c \in \mathbb{R}^d$ and imagine growing B by increasing its radius starting from zero. For nearly all starting points c, the ball B eventually intersects one, *and only one*, of the $\mathcal{M}_i$'s. Thus the nearest point to c on $\mathcal{M}$, in the norm L_p, lies on $\mathcal{M}_i$. (Note that the nearest point on $\mathcal{M}_i$ need not be unique.)

The *decision axis* Λ_p of $\mathcal{M}$ is the set of points c such that the boundary of B intersects two or more of the $\mathcal{M}_i$, but the interior of B does not intersect $\mathcal{M}$ at all. In other words, the decision axis Λ_p is the set of points that have two or more closest points, in the norm L_p, *on distinct class manifolds*. See Figure 4.1. The decision axis is inspired by the medial axis, which was first proposed by Blum [9] in the context of image analysis and subsequently modified for the purposes of curve and surface reconstruction by Amenta et al. [3, 6]. We have modified the definition to account for multiple class manifolds and have renamed our variant in order to avoid confusion in the future.

The decision axis Λ_p can intuitively be thought of as a decision boundary that is optimal in the following sense. First, Λ_p separates the class manifolds when they do not intersect. Second, each point of Λ_p is as far away from the class manifolds as possible in the norm L_p. As shown in the leftmost example in Figure 4.1, in the case of two linearly separable circles of equal radius, the decision axis Λ_2 is exactly the line that separates the data with maximum margin. For arbitrary manifolds, Λ_p generalizes the notion of maximum margin to account for the arbitrary curvature of the class manifolds.

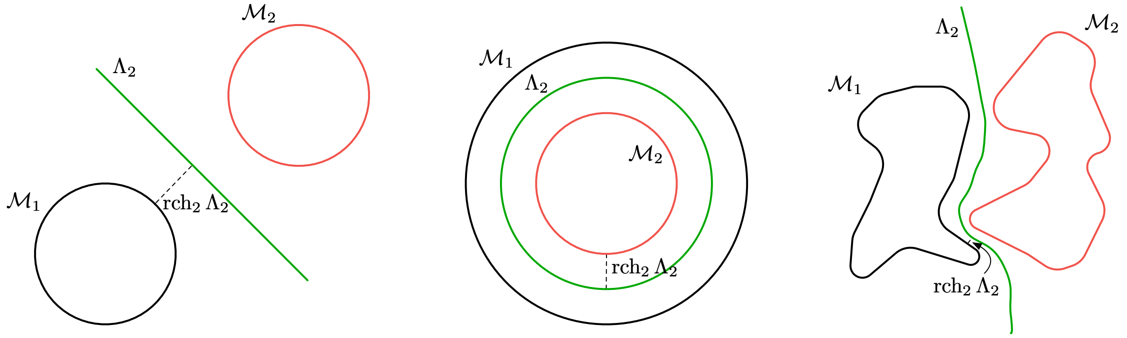

Figure 4.1: Examples of the decision axis Λ_2, shown here in green, for different data manifolds. Intuitively, the decision axis captures an optimal decision boundary between the data manifolds. It's optimal in the sense that each point on the decision axis is as far away from each data manifold as possible. Notice that in the first example, the decision axis coincides with the maximum margin line.

Let $T \subset \mathbb{R}^d$ be any set. The *reach* $\text{rch}_p\,(T; \mathcal{M})$ of $\mathcal{M}$ is defined as $\inf_{x \in \mathcal{M}, y \in T} \|x - y\|_p$. When $\mathcal{M}$ is compact, the reach is achieved by the point on $\mathcal{M}$ that is closest to T under the L_p norm. We will drop $\mathcal{M}$ from the notation when it is understood from context.

Finally, an ϵ-*tubular neighborhood* of $\mathcal{M}$ is defined as $\mathcal{M}^{\epsilon,p} = \{x \in \mathbb{R}^d : \inf_{y \in \mathcal{M}} \|x - y\|_p \leq \epsilon\}$. That is, $\mathcal{M}^{\epsilon,p}$ is the set of all points whose distance to $\mathcal{M}$ under the metric induced by L_p is less than ϵ. Note that while $\mathcal{M}$ is k-dimensional, $\mathcal{M}^{\epsilon,p}$ is always d-dimensional. Tubular neighborhoods are how we rigorously define adversarial examples. Consider a classifier $f : \mathbb{R}^d \to [C]$ for $\mathcal{M}$. An ϵ-*adversarial example* is a point $x \in \mathcal{M}_i^{\epsilon,p}$ such that $f(x) \neq i$. A classifier f is robust to all ϵ-adversarial examples when f correctly classifies not only $\mathcal{M}$, but all of $\mathcal{M}^{\epsilon,p}$. Thus the problem of being robust to adversarial examples is rightly seen as one of *generalization*. In this chapter we will be primarily concerned with exploring the conditions under which we can provably learn a decision boundary that correctly classifies $\mathcal{M}^{\epsilon,p}$. When $\epsilon < \text{rch}_p\,\Lambda_p$, the decision axis Λ_p is one decision boundary that correctly classifies $\mathcal{M}^{\epsilon,p}$. Throughout the remainder of this chapter we will drop the p in $\mathcal{M}^{\epsilon,p}$ from the notation, instead writing $\mathcal{M}^{\epsilon}$; the norm will always be clear from context.

The geometric quantities defined above can be defined more generally for any distance metric $d(\cdot,\cdot)$. In this chapter we will focus exclusively on the metrics induced by the norms L_p for $p > 0$. The decision axis under L_2 is in general *not* identical to the decision axis under L_∞. In Section 4.4 we will prove that since Λ_2 is not identical to Λ_∞ there exists a tradeoff in the robustness of any decision boundary between the two norms.

4.4 A Provable Tradeoff in Robustness Between Norms

Schott et al. [78] explore the vulnerability of robust classifiers to attacks under different norms. In particular, they take the robust pretrained classifier of [66], which was trained to be robust to L_∞-perturbations, and subject it to L_0 and L_2 attacks. They show that accuracy drops to 0% under L_0 attacks and to 35% under L_2. Here we explain why poor robustness under the norm L_2 should be expected.

We say a decision boundary $\mathcal{D}_f$ for a classifier f is ϵ-robust in the L_p norm if $\epsilon < \mathrm{rch}_p\,\mathcal{D}_f$. In words, starting from any point $x \in \mathcal{M}$, a perturbation η_x must have p-norm greater than $\mathrm{rch}_p\,\mathcal{D}_f$ to cross the decision boundary. The most robust decision boundary to L_p-perturbations is Λ_p. In Theorem 24 we construct a learning setting where Λ_2 is distinct from Λ_∞. Thus, in general, *no single decision boundary can be optimally robust in all norms.*

Theorem 24. *Let $S_1, S_2 \subset \mathbb{R}^{d+1}$ be two concentric d-spheres with radii $r_1 < r_2$ respectively. Let $S = S_1 \cup S_2$ and let $\Lambda_2, \Lambda_\infty$ be the L_2 and L_∞ decision axes of S. Then $\Lambda_2 \neq \Lambda_\infty$. Furthermore* $\mathrm{rch}_2\,\Lambda_\infty \in O(\mathrm{rch}_2\,\Lambda_2/\sqrt{d})$.

Proof. The decision axis under L_2, Λ_2, is just the d-sphere with radius $(r_1 + r_2)/2$. However, Λ_∞ is *not* identical to Λ_2 in this setting; in fact most Λ_∞ of approaches S_1 as d increases.

The geometry of a L_∞-ball B_Δ centered at $m \in \mathbb{R}^d$ with radius Δ is that of a hypercube centered at m with side length 2Δ. To find a point on Λ_∞ we place B_Δ tangent to the north pole q of S_1 so that the corners of B_Δ touch S_2. The north pole has coordinate representation $q = (0, \ldots, 0, r_1)$, the center $m = (0, \ldots, 0, r_1 + \Delta)$, and a corner of B_Δ can be expressed as $p = (\Delta, \ldots, \Delta, r_1 + 2\Delta)$. Additionally we have the constraint that $\|p\|_2 = r_2$ since $p \in S_2$. Then we can solve for Δ as

$$r_2^2 = \|p\|_2^2 = (d-1)\Delta^2 + (r_1 + 2\Delta)^2 = (d+3)\Delta^2 + 4r_1\Delta + r_1^2;$$
$$\Delta = \frac{-2r_1 + \sqrt{r_1^2 + 3r_2^2 + d(r_2^2 - r_1^2)}}{d+3},$$

where the last step follows from the quadratic formula and the fact that $\Delta > 0$. For fixed r_1, r_2, the value Δ scales as $O(1/\sqrt{d})$. It follows that $\mathrm{rch}_2\,\Lambda_\infty \in O(\mathrm{rch}_2\,\Lambda_2/\sqrt{d})$. □

From Theorem 24 we conclude that the minimum distance from S_1 to Λ_∞ *under the L_2 norm* is upper bounded as $\mathrm{rch}_2\,\Lambda_\infty \in O(\mathrm{rch}_2\,\Lambda_2/\sqrt{d})$. If a classifier f is trained to learn Λ_∞, an adversary,

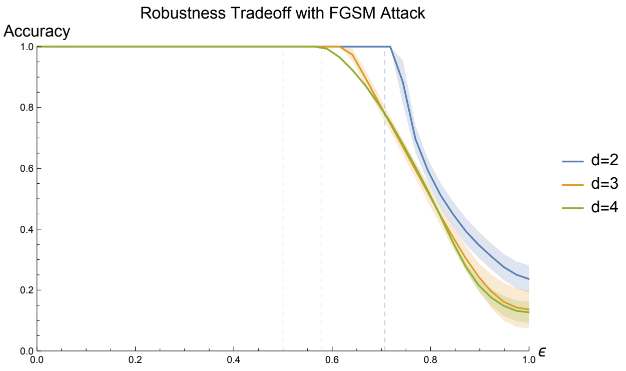

Figure 4.2: As the dimension increases, the $\mathrm{rch}_2\,(\Lambda_\infty; S_1 \cup S_2)$ decreases, and so an L_∞ robust classifier is less robust to L_2 attacks. The dashed lines are placed at $1/\sqrt{d}$, where our theoretical results suggest we should start finding L_2 adversarial examples. We use the robust L_∞ loss of [94]

starting on S_1, can construct an L_2 adversarial example for a perturbation as small as $O(1/\sqrt{d})$. Thus we should *expect* f to be less robust to L_2-perturbations. Figure 4.2 verifies this result experimentally.

We expect that $\Lambda_2 \neq \Lambda_\infty$ is the common case in practice. For example, Theorem 24 extends immediately to concentric cylinders and intertwined tori by considering 2-dimensional planar cross-sections. In general, we expect that $\Lambda_2 \neq \Lambda_\infty$ in situations where a 2-dimensional cross-section with $\mathcal{M}$ has nontrivial curvature.

Theorem 24 is important because, even in recent literature, researchers have attributed this phenomena to overfitting. Schott et al. [78] state that "the widely recognized and by far most successful defense by Madry et al. (1) *overfits* on the L_∞ metric (it's highly susceptible to L_2 and L_0 perturbations)" (emphasis ours). We disagree; the Madry et al. [66] classifier performed exactly as intended. It learned a decision boundary that is robust under L_∞, which we have shown is quite different from the most robust decision boundary under L_2.

Interestingly, the proposed models of Schott et al. [78] also suffer from this tradeoff. Their model ABS has accuracy 80% to L_2 attacks but drops to 8% for L_∞. Similarly their model ABS Binary has accuracy 77% to L_∞ attacks but drops to 39% for L_2 attacks.

We reiterate, in general, no single decision boundary can be optimally robust in all norms.

4.5 Provably Robust Classifiers

Adversarial training, the process of training on adversarial examples generated in a L_p-ball around the training data, is a very natural approach to constructing robust models [41, 66]. In our notation this corresponds to training on samples drawn from X^ϵ for some ϵ. While natural, we show that there are simple settings where this approach is much less sample-efficient than other classification algorithms, if the *only* guarantee is correctness in X^ϵ.

Define a learning algorithm $\mathcal{L}$ with the property that, given a training set $X \subset \mathcal{M}$ sampled from a manifold $\mathcal{M}$, $\mathcal{L}$ outputs a model $f_{\mathcal{L}}$ such that for every $x \in X$ with label y, and every $\hat{x} \in B(x, \operatorname{rch}_p \Lambda_p)$, $f_{\mathcal{L}}(\hat{x}) = f_{\mathcal{L}}(x) = y$. Here $B(x, r)$ denotes the ball centered at x of radius r in the relevant norm. That is, $\mathcal{L}$ learns a model that outputs the same label for any L_p-perturbation of x up to $\operatorname{rch}_p \Lambda_p$ as it outputs for x. $\mathcal{L}$ is our theoretical model of adversarial training [41, 66]. Theorem 25 states that $\mathcal{L}$ is sample inefficient in high codimensions.

Theorem 25. *There exists a classification algorithm $\mathcal{A}$ that, for a particular choice of $\mathcal{M}$, correctly classifies $\mathcal{M}^\epsilon$ using exponentially fewer samples than are required for $\mathcal{L}$ to correctly classify $\mathcal{M}^\epsilon$.*

Theorem 25 follows from Theorems 26 and 27. In Theorems 26 and 27 we will prove that a nearest neighbor classifier f_{nn} is one such classification algorithm. Nearest neighbor classifiers are naturally robust in high codimensions because the Voronoi cells of X are *elongated in the directions normal* to $\mathcal{M}$ when X is dense [28].

Before we state Theorem 26 we must introduce a sampling condition on $\mathcal{M}$. A δ-cover of a manifold $\mathcal{M}$ in the norm L_p is a finite set of points X such that for every $x \in \mathcal{M}$ there exists X_i such that $\|x - X_i\|_p \leq \delta$. Theorem 26 gives a sufficient sampling condition for $f_{\mathcal{L}}$ to correctly classify $\mathcal{M}^\epsilon$ for all manifolds $\mathcal{M}$. Theorem 26 also provides a sufficient sampling condition for a nearest neighbor classifier f_{nn} to correctly classify $\mathcal{M}^\epsilon$, which is substantially less dense than that of $f_{\mathcal{L}}$. Thus different classification algorithms have different sampling requirements in high codimensions.

Theorem 26. *Let $\mathcal{M} \subset \mathbb{R}^d$ be a k-dimensional manifold and let $\epsilon < \operatorname{rch}_p \Lambda_p$ for any $p > 0$. Let f_{nn} be a nearest neighbor classifier and let $f_{\mathcal{L}}$ be the output of a learning algorithm $\mathcal{L}$ as described above. Let $X_{\text{nn}}, X_{\mathcal{L}} \subset \mathcal{M}$ denote the training sets for f_{nn} and $\mathcal{L}$ respectively. We have the following sampling guarantees:*

1. *If X_{nn} is a δ-cover for $\delta \leq 2(\operatorname{rch}_p \Lambda_p - \epsilon)$ then f_{nn} correctly classifies $\mathcal{M}^\epsilon$.*
2. *If $X_{\mathcal{L}}$ is a δ-cover for $\delta \leq \operatorname{rch}_p \Lambda_p - \epsilon$ then $f_{\mathcal{L}}$ correctly classifies $\mathcal{M}^\epsilon$.*

Proof. Here we use $d(\cdot, \cdot)$ to denote the metric induced by the L_p norm. We begin by proving (1). Let $q \in \mathcal{M}^\epsilon$ be any point in $\mathcal{M}^\epsilon$. Suppose without loss of generality that $q \in \mathcal{M}_i^\epsilon$ for some class i. The distance $d(q, \mathcal{M}_j)$ from q to any other data manifold $\mathcal{M}_j$, and thus any sample on $\mathcal{M}_j$, is lower bounded by $d(q, \mathcal{M}_j) \geq 2 \operatorname{rch}_p \Lambda_p - \epsilon$. See Figure 4.3. It is then both necessary and sufficient that there exists a $x \in \mathcal{M}_i$ such that $d(q, x) < 2 \operatorname{rch}_p \Lambda_p - \epsilon$ for $f_{\text{nn}}(q) = i$. (Necessary since a properly placed sample on $\mathcal{M}_j$ can achieve the lower bound on $d(q, \mathcal{M}_j)$.) The distance from q to the nearest sample x on $\mathcal{M}_i$ is $d(q, x) \leq \epsilon + \delta$ for some $\delta > 0$. The question is how large can we allow δ to be and still guarantee that f_{nn} correctly classifies $\mathcal{M}^\epsilon$? We need

$$d(q, x) \leq \epsilon + \delta \leq 2 \operatorname{rch}_p \Lambda_p - \epsilon \leq d(q, \mathcal{M}_j)$$

which implies that $\delta \leq 2(\operatorname{rch}_p \Lambda_p - \epsilon)$. It follows that a δ-cover with $\delta = 2(\operatorname{rch}_p \Lambda_p - \epsilon)$ is sufficient, and in some cases necessary, to guarantee that f_{nn} correctly classifies $\mathcal{M}^\epsilon$.

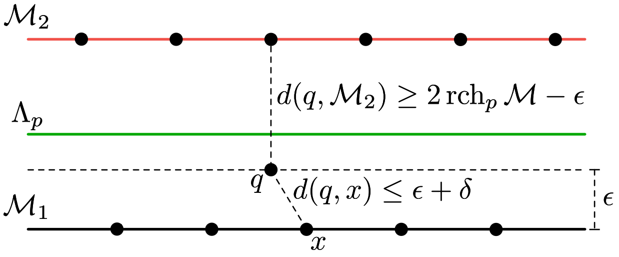
$\mathcal{M}_2$
$d(q, \mathcal{M}_2) \geq 2\,\mathrm{rch}_p\,\mathcal{M} - \epsilon$
Λ_p
q
$d(q, x) \leq \epsilon + \delta$
ϵ
$\mathcal{M}_1$
x

Figure 4.3: Proof of Theorem 26. The distance from a query point q to $\mathcal{M}_2$, and thus the closest incorrectly labeled sample, is lower bounded by the distance necessary to reach the medial axis Λ_p plus the distance from Λ_p to $\mathcal{M}_2$.

Next we prove (2). As before let $q \in \mathcal{M}_i^\epsilon$. It is both necessary and sufficient for $q \in B(x, \mathrm{rch}_p\,\Lambda_p)$ for some sample $x \in \mathcal{M}_i$ to guarantee that $f_{\mathcal{L}}(q) = i$, by definition of $\mathcal{L}$. The distance to the nearest sample x on $\mathcal{M}_i$ is $d(q, x) \leq \epsilon + \delta$ for some $\delta > 0$. Thus it suffices that $\delta \leq \mathrm{rch}_p\,\Lambda_p - \epsilon$. □

The bounds on δ in Theorem 26 are sufficient, but they are not always necessary. There exist manifolds where the bounds in Theorem 26 are pessimistic, and less dense samples corresponding to larger values of δ would suffice.

Next we will show a setting where bounds on δ similar to those in Theorem 26 are *necessary*. In this setting, the difference of a factor of 2 in δ between the sampling requirements of f_{nn} and $f_{\mathcal{L}}$ leads to an exponential gap between the sizes of X_{nn} and $X_{\mathcal{L}}$ necessary to achieve the same amount of robustness.

Define $\Pi_1 = \{x \in \mathbb{R}^d : \ell \leq x_1, \ldots, x_k \leq \mu \text{ and } x_{k+1} = \ldots = x_d = 0\}$; that is Π_1 is a subset of the x_1-...-x_k-plane bounded between the coordinates $[\ell, \mu]$. Similarly define $\Pi_2 = \{x \in \mathbb{R}^d : \ell \leq x_1, \ldots, x_k \leq \mu \text{ and } x_{k+1} = \ldots = x_{d-1} = 0 \text{ and } x_d = 2\}$. Note that Π_2 lies in the subspace $x_d = 2$; thus $\mathrm{rch}_2\,\Lambda_2 = 1$, where Λ_2 is the decision axis of $\Pi = \Pi_1 \cup \Pi_2$. In the L_2 norm we can show that the gap in Theorem 26 is necessary for $\Pi = \Pi_1 \cup \Pi_2$. Furthermore the bounds we derive for δ-covers for Π for both f_{nn} and $f_{\mathcal{L}}$ are tight. Combined with well-known properties of covers, we get that the ratio $|X_{\mathcal{L}}|/|X_{\mathrm{nn}}|$ is exponential in k.

Theorem 27. *Let $\Pi = \Pi_1 \cup \Pi_2$ as described above. Let $X_{\mathrm{nn}}, X_{\mathcal{L}} \subset \Pi$ be minimum training sets necessary to guarantee that f_{nn} and $f_{\mathcal{L}}$ correctly classify $\mathcal{M}^\epsilon$. Then we have that*

$$\frac{|X_{\mathcal{L}}|}{|X_{\mathrm{nn}}|} \geq 2^{k/2} \tag{4.1}$$

Proof. Let $q \in \Pi_1^\epsilon$. Since Π_1 is flat, the distance from q to the nearest sample $x \in \Pi_1$ is bounded as $\|q - x\|_2 \leq \sqrt{\epsilon^2 + \delta^2}$. For $f_{\mathrm{nn}}(q) = 1$ we need that $\|q - x\|_2 \leq 2 - \epsilon$, and so it suffices that $\delta \leq 2\sqrt{1 - \epsilon}$. In this setting, this is also necessary; should δ be any larger a property placed sample on Π_2 can claim q in its Voronoi cell.

Similarly for $f_{\mathcal{L}}(q) = 1$ we need that $\|q - x\|_2 \leq 1$, and so it suffices that $\delta \leq \sqrt{1-\epsilon^2}$. In this setting, this is also necessary; should δ be any larger, q lies outside of every L_2-ball $B(x, 1)$ and so $\mathcal{L}$ is free to learn a decision boundary that misclassifies q.

Let $\mathcal{N}(\delta, \mathcal{M})$ denote the size of the minimum δ-cover of $\mathcal{M}$. Since Π is flat (has no curvature) and since the intersection of Π with a d-ball centered at a point on Π is a k-ball, a standard volume argument can be applied in the affine subspace aff Π to conclude that $\mathcal{N}(\delta, \Pi) \in \Theta\left(\mathrm{vol}_k \Pi/\delta^k\right)$. So we have

$$\begin{aligned}\frac{\mathcal{N}(\sqrt{1-\epsilon^2}, \Pi)}{\mathcal{N}(2\sqrt{1-\epsilon}, \Pi)} &= 2^k\left(\frac{1}{1+\epsilon}\right)^{k/2}\\ &\geq 2^{k/2}\end{aligned}$$

Since Π is constant in both settings, the factor $\mathrm{vol}_k \Pi$ as well as the constant factors hidden by $\Theta(\cdot)$ cancel. (Note that we are using the fact that Π_1, Π_2 have finite k-dimensional volume.) The inequality follows from the fact that the expression $(1+\epsilon)^{-k/2}$ is monotonically decreasing on the interval $[0, 1]$ and takes value $2^{-k/2}$ at $\epsilon = 1$. □

We have shown that both $\mathcal{L}$ and nearest neighbor classifiers learn robust decision boundaries when provided sufficiently dense samples of $\mathcal{M}$. However there are settings where nearest neighbors is exponentially more sample-efficient than $\mathcal{L}$ in achieving the same amount of robustness. We experimentally verify these theoretical results in Section 4.10.

4.6 X^ϵ is a Poor Model of $\mathcal{M}^\epsilon$

Madry et al. [66] suggest training a robust classifier with the help of an adversary which, at each iteration, produces ϵ-perturbations around the training set that are incorrectly classified. In our notation, this corresponds to learning a decision boundary that correctly classifies $X^\epsilon = \{x \in \mathbb{R}^d : \|x - X_i\|_2 \leq \epsilon$ for some training point $X_i\}$. We believe this approach is insufficiently robust in practice, as X^ϵ is often a poor model for $\mathcal{M}^\epsilon$. In this section, we show that the volume $\mathrm{vol}\, X^\epsilon$ is often a vanishingly small percentage of $\mathrm{vol}\,\mathcal{M}^\epsilon$. These results shed light on why the ball-based learning algorithm $\mathcal{L}$ defined in Section 4.5 is so much less sample-efficient than nearest neighbor classifiers. In Section 4.10 we experimentally verify these observations by showing that in high-dimensional space it is easy to find adversarial examples even after training against a strong adversary. For the remainder of this section we will consider the L_2 norm.

Theorem 28. *Let $\mathcal{M} \subset \mathbb{R}^d$ be a k-dimensional manifold embedded in $\mathbb{R}^d$ such that $\mathrm{vol}_k \mathcal{M} < \infty$. Let $X \subset \mathcal{M}$ be a finite set of points sampled from $\mathcal{M}$. Suppose that $\epsilon \leq \mathrm{rch}_2\, \Xi$ where Ξ is the medial axis of $\mathcal{M}$, defined as in [28]. Then the percentage of $\mathcal{M}^\epsilon$ covered by X^ϵ is upper bounded by*

$$\frac{\mathrm{vol}\, X^\epsilon}{\mathrm{vol}\,\mathcal{M}^\epsilon} \leq \frac{\pi^{k/2}\Gamma(\frac{d-k}{2}+1)}{\Gamma(\frac{d}{2}+1)}\frac{\epsilon^k}{\mathrm{vol}_k \mathcal{M}}|X| \in O\left(\left(\frac{2\pi}{d-k}\right)^{k/2}\frac{\epsilon^k}{\mathrm{vol}_k \mathcal{M}}|X|\right). \tag{4.2}$$

As the codimension $(d-k) \to \infty$, Equation 4.2 approaches 0, *for any fixed* $|X|$.

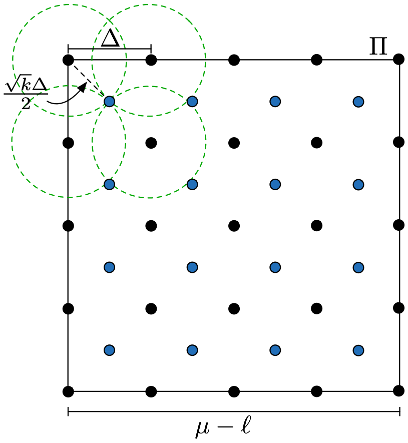

Figure 4.4: To construct an δ-cover we place sample points, shown here in black, along a regular grid with spacing Δ. The blue points are the furthest points of Π from the sample. To cover Π we need $\Delta = 2\delta/\sqrt{k}$.

Proof. Assuming the balls centered on the samples in X are disjoint we get the upper bound

$$\operatorname{vol} X^{\epsilon} \leq \operatorname{vol} B_{\epsilon}|X| = \frac{\pi^{d/2}}{\Gamma(\frac{d}{2}+1)}\epsilon^{d}|X|. \tag{4.3}$$

This is identical to the reasoning in Equation 4.5.

The medial axis Ξ of $\mathcal{M}$ is defined as the closure of the set of all points in $\mathbb{R}^d$ that have two or more closest points on $\mathcal{M}$ in the norm L_2. The medial axis Ξ is similar to the decision axis Λ_2, except that the nearest points do not need to be on distinct class manifolds. For $\epsilon \leq \operatorname{rch}_2 \Xi$, we have the lower bound

$$\operatorname{vol} \mathcal{M}^{\epsilon} \geq \operatorname{vol}_{d-k} B_{\epsilon}^{d-k} \operatorname{vol}_k \mathcal{M} = \frac{\pi^{(d-k)/2}}{\Gamma\left(\frac{d-k}{2}+1\right)}\epsilon^{d-k} \operatorname{vol}_k \mathcal{M}. \tag{4.4}$$

Combining Equations 4.3 and 4.4 gives the result. To get the asymptotic result we apply

Stirling's approximation to get

$$\begin{aligned}\frac{\Gamma(\frac{d-k}{2}+1)}{\Gamma(\frac{d}{2}+1)} &\approx (2e)^{k/2}\frac{(d-k)^{(d-k+1)/2}}{d^{(d+1)/2}}\\ &= (2e)^{k/2}\frac{\left(\frac{d-k}{d}\right)^{(d+1)/2}}{(d-k)^{k/2}}\\ &= (2e)^{k/2}\frac{\left(1-\frac{k}{d}\right)^{(d+1)/2}}{(d-k)^{k/2}}\\ &\approx \left(\frac{2}{d-k}\right)^{k/2}.\end{aligned}$$

The last step follows from the fact that $\lim_{d\to\infty}(1-k/d)^{(d+1)/2} = e^{-k/2}$, where e is the base of the natural logarithm. □

In high codimension, even moderate under-sampling of $\mathcal{M}$ leads to a significant loss of coverage of $\mathcal{M}^\epsilon$ because the volume of the union of balls centered at the samples shrinks faster than the volume of $\mathcal{M}^\epsilon$. Theorem 28 states that in high codimensions the fraction of $\mathcal{M}^\epsilon$ covered by X^ϵ goes to 0. Almost nothing is covered by X^ϵ for training set sizes that are realistic in practice. Thus X^ϵ is a poor model of $\mathcal{M}^\epsilon$, and high classificaiton accuracy on X^ϵ does not imply high accuracy in $\mathcal{M}^\epsilon$.

Note that an alternative way of defining the ratio $\operatorname{vol} X^\epsilon / \operatorname{vol} \mathcal{M}^\epsilon$ is as $\operatorname{vol}(X^\epsilon \cap \mathcal{M}^\epsilon)/\operatorname{vol}\mathcal{M}^\epsilon$. This is equivalent in our setting since $X \subset \mathcal{M}$ and so $X^\epsilon \subset \mathcal{M}^\epsilon$.

For the remainder of the section we provide intuition for Theorem 28 by considering the special case of k-dimensional planes. Define $\Pi = \{x \in \mathbb{R}^d : \ell \leq x_1, \ldots, x_k \leq \mu \text{ and } x_{k+1} = \ldots = x_d = 0\}$; that is Π is a subset of the x_1-. . .-x_k-plane bounded between the coordinates $[\ell, \mu]$. Recall that a δ-cover of a manifold $\mathcal{M}$ in the norm $\|\cdot\|_2$ is a finite set of points X such that for every $x \in \mathcal{M}$ there exists X_i such that $\|x - X_i\|_2 \leq \delta$. It is easy to construct an *explicit* δ-cover X of Π: place sample points at the vertices of a regular grid, shown in Figure 4.4 by the black vertices. The centers of the cubes of this regular grid, shown in blue in Figure 4.4, are the furthest points from the samples. The distance from the vertices of the grid to the centers is $\sqrt{k}\Delta/2$ where Δ is the spacing between points along an axis of the grid. To construct a δ-cover we need $\sqrt{k}\Delta/2 = \delta$ which gives a spacing of $\Delta = 2\delta/\sqrt{k}$. The size of this sample is $|X| = \left(\frac{\sqrt{k}(\mu-\ell)}{2\delta}\right)^k$. Note that $|X|$ scales exponentially in k, the dimension of Π, not in d, the dimension of the embedding space.

Recall that Π^δ is the δ-tubular neighborhood of Π. The δ-balls around X, which comprise X^δ, cover Π and so any robust approach that guarantees correct classification within X^δ will achieve perfect accuracy on Π. However, we will show that X^δ covers only a vanishingly small fraction of Π^δ. Let B_δ denote the d-ball of radius δ centered at the origin. An upper bound on the volume of X^δ is

$$\operatorname{vol} X^\delta \leq \operatorname{vol} B_\delta |X| = \frac{\pi^{d/2}}{\Gamma(\frac{d}{2}+1)}\delta^d\left(\frac{\sqrt{k}(\mu-\ell)}{2\delta}\right)^k = \frac{\pi^{d/2}}{\Gamma(\frac{d}{2}+1)}\delta^{(d-k)}\left(\frac{\sqrt{k}(\mu-\ell)}{2}\right)^k. \tag{4.5}$$

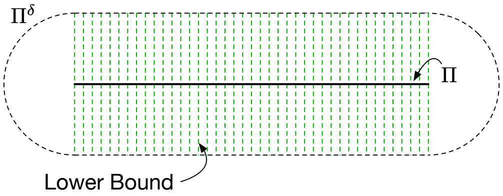

Figure 4.5: An illustration of the lower bound technique used in Equation 4.6. The volume $\operatorname{vol}\Pi^\delta$ shown in the black dashed lines, is bounded from below by placing a $(d-k)$-dimensional ball of radius δ at each point of Π, shown in green. In this illustration, a 1-dimensional manifold is embedded in 2 dimensions, so these balls are 1-dimensional line segments.

Next we bound the volume $\operatorname{vol}\Pi^\delta$ from below. Intuitively, a lower bound on the volume can be derived by placing a $(d-k)$-dimensional ball in the normal space at each point of Π and integrating the volumes. Figure 4.4 (Right) illustrates the lower bound argument in the case of $k = 1, d = 2$.

$$\operatorname{vol}\Pi^\delta \geq \operatorname{vol}_{d-k} B_\delta^{d-k} \operatorname{vol}_k \Pi = \frac{\pi^{(d-k)/2}}{\Gamma\left(\frac{d-k}{2}+1\right)}\delta^{d-k}(\mu-\ell)^k. \tag{4.6}$$

Combining Equations 4.5 and 4.6 gives an upper bound on the percentage of Π^δ that is covered by X^ϵ.

$$\frac{\operatorname{vol}X^\delta}{\operatorname{vol}\Pi^\delta} \leq \frac{\pi^{k/2}\Gamma\left(\frac{d-k}{2}+1\right)}{\Gamma\left(\frac{d}{2}+1\right)}\left(\frac{\sqrt{k}}{2}\right)^k. \tag{4.7}$$

Notice that the factors involving δ and $(\mu-\ell)$ cancel. Figure 4.6 (Left) shows that this expression approaches 0 as the codimension $(d-k)$ of Π increases.

Suppose we set $\delta = 1$ and construct a 1-cover of Π. The number of points necessary to cover Π with balls of radius 1 depends *only* on k, not the embedding dimension d. However the number of points necessary to cover the tubular neighborhood Π^1 with balls of radius 1 increases depends on *both k* and d. In Theorem 29 we derive a lower bound on the number of samples necessary to cover Π^1.

Theorem 29. *Let Π be a bounded k-flat as described above, bounded along each axis by $\ell < \mu$. Let n denote the number of samples necessary to cover the* 1*-tubular neighborhood Π^1 of Π with L_2-balls of radius* 1*. That is let n be the minimum value for which there exists a finite sample X of size n such that $\Pi^1 \subset \cup_{x\in X}B(x,1) = X^1$. Then*

$$n \geq \frac{\pi^{-k/2}\Gamma\left(\frac{d}{2}+1\right)}{\Gamma\left(\frac{d-k}{2}+1\right)}(\mu-\ell)^k \in \Omega\left(\left(\frac{d-k}{2\pi}\right)^{k/2}(\mu-\ell)^k\right). \tag{4.8}$$

Proof. We first construct an upper bound by generously assuming that the balls centered at the samples are disjoint. That is

$$\frac{\operatorname{vol} X^\delta}{\operatorname{vol} \Pi^\delta} \leq \frac{n \operatorname{vol} B_\delta}{\operatorname{vol} \Pi^\delta}. \tag{4.9}$$

To guarantee that $\Pi^1 \subset \cup_{x \in X} B(x, 1) = X^1$ we set the left hand side of Equation 4.9 equal to 1 and solve for n.

$$\begin{aligned} 1 = \frac{\operatorname{vol} X^\delta}{\operatorname{vol} \Pi^\delta} &\leq \frac{n \operatorname{vol} B_\delta}{\operatorname{vol} \Pi^\delta} \\ n &\geq \frac{\operatorname{vol} \Pi^\delta}{\operatorname{vol} B_\delta} \\ &\geq \frac{\pi^{-k/2} \Gamma\left(\frac{d}{2} + 1\right)}{\Gamma\left(\frac{d-k}{2} + 1\right)} \left(\frac{\mu - \ell}{\delta}\right)^k \end{aligned}$$

The last inequality follows from Equation 4.6. Setting $\delta = 1$ gives the result. The asymptotic result is similar to the argument in the proof of Theorem 28. □

Theorem 29 states that, in general, it takes many fewer samples to accurately model $\mathcal{M}$ than to model $\mathcal{M}^\epsilon$. Figure 4.6 (Right) compares the number of points necessary to construct a 1-cover of Π with the lower bound on the number necessary to cover Π^1 from Theorem 29. The number of points necessary to cover Π^1 increases as $\Omega\left((d - k)^{k/2}\right)$, scaling polynomially in d and exponentially in k. In contrast, the number necessary to construct a 1-cover of Π remains constant as d increases, depending only on k.

Our lower bound of $\Omega\left((d - k)^{k/2}\right)$ samples is similar to the work of [76] who prove that, in the simple Gaussian setting, robustness *requires* as much as $\Omega(\sqrt{d})$ more samples. Their arguments are statistical while ours are geometric.

Approaches that produce robust classifiers by generating adversarial examples in the ϵ-balls centered on the training set do not accurately model $\mathcal{M}^\epsilon$, and it will take *many* more samples to do so. If the method behaves arbitrarily outside of the ϵ-balls that define X^ϵ, adversarial examples will still exist and it will likely be easy to find them. The reason deep learning has performed so well on a variety of tasks, in spite of the brittleness made apparent by adversarial examples, is because it is much easier to perform well on $\mathcal{M}$ than it is to perform well on $\mathcal{M}^\epsilon$.

4.7 A Lower Bound on Model Expressiveness

A Simple Example

Consider the case of two concentric circles C_1, C_2 with radii $r_1 < r_2$ respectively, as illustrated in Figure 4.7. Each circle represents a different class of data. Suppose that we train a parametric model $f(x; \boldsymbol{\theta})$ with p parameters so that for $x \in C_1$, $f(x; \boldsymbol{\theta}) > 0$ and for $x \in C_2$, $f(x; \boldsymbol{\theta}) < 0$. How does the

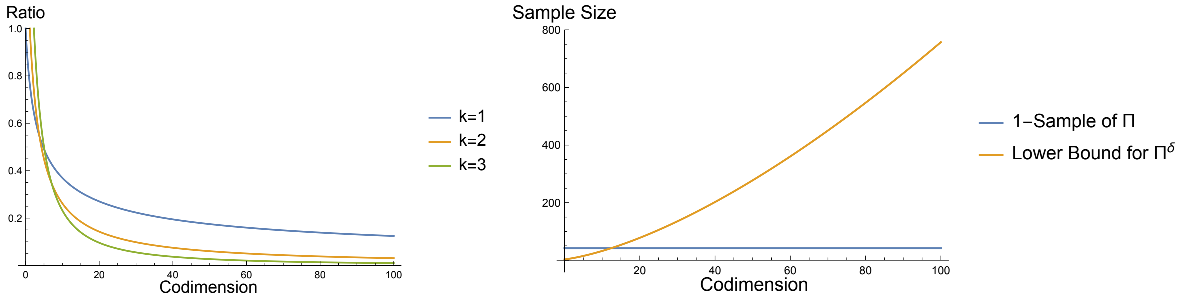

Figure 4.6: We plot the upper bound in Equation 4.7 on the left. As the codimension increases, the percentage of volume of Π^1 covered by 1-balls around the 1-sample approaches 0. On the right we plot the number of samples necessary to cover Π, shown in blue, against the number of samples necessary to cover Π^1, shown in orange, as the codimension increases.

number of parameters p necessary to ensure that such a decision boundary can be expressed by $f(\cdot;\boldsymbol{\theta})$ increase as the gap between C_1 and C_2 decreases?

Suppose that we first lift C_1 and C_2 to a parabola in $\mathbb{R}^3$ via map $\phi(x_1, x_2) = (x_1, x_2, x_1^2 + x_2^2)$. That is, we construct the sets $C_1^+ = \{\phi(x_1, x_2) : (x_1, x_2) \in C_1\}$ and similarly for C_2^+. After applying ϕ, C_1^+ and C_2^+ are *linearly separable* for any $r_2 - r_1 > 0$. The linear decision boundary in $\mathbb{R}^3$ maps back to a circle in $\mathbb{R}^2$ that separates C_1 and C_2. This is not the case for deep networks; the number of parameters necessary to separate C_1 and C_2 will depend on the gap $r_2 - r_1$.

In the important special case where f is parameterized by a fully connected deep network with ℓ layers, h hidden units per layer, and ReLU activations, Raghu et al. [71] prove that f subdivides the input space into convex polytopes. In each convex polytope, f defines a linear function that agrees on the boundary of the polytope with its neighbors. They showed that, when the inputs are in $\mathbb{R}^2$, the number of polytopes in the subdivision is at most $\mathcal{O}(h^{2\ell})$ [71][Theorem 1].

Let $\mathcal{S}_f$ denote the subdivision of space into convex polytopes induced by f. Consider the decision boundary $\mathcal{D}_f = \{x \in \mathbb{R}^d : f(x;\boldsymbol{\theta}) = 0\}$ of f. $\mathcal{D}_f$ can be constructed by examining each polytope $P \in \mathcal{S}_f$ and solving the linear equation $f_P(x) = 0$ where f_P is the linear function defined on P by f. Since f_P is linear the solution is either (1) the empty set, (2) a *single* line segment, or (3) all of P. Case (3) is a degenerate case and there are ways to perturb f by an infinitesimally small amount such that case (3) never occurs and the classification accuracy is unchanged. Thus we conclude that $\mathcal{D}_f$ is a piecewise-linear curve comprised of line segments. (In higher dimensions $\mathcal{D}_f$ is composed of subsets of hyperplanes.) See Figure 4.7.

Suppose that $\mathcal{D}_f$ separates C_1 from C_2 and let $s \in \mathcal{D}_f$ be a line segment of the decision boundary. Since s lies in the space between C_1 and C_2, the length $|s| \leq 2\sqrt{r_2^2 - r_1^2}$, which is tight when s is tangent to C_1 and touches C_2 at both of its endpoints. For $\mathcal{D}_f$ to separate C_1 from C_2, $\mathcal{D}_f$ must make a full rotation of 2π around the origin. The portion of this rotation that s can contribute is upper bounded by $2 \arccos \frac{r_1}{r_2}$. Thus the number of line segments that comprise $\mathcal{D}_f$ is lower bounded by $\frac{\pi}{\arccos \frac{r_1}{r_2}}$.

As $r_2 \to r_1$, the minimum number of segment necessary to separate C_1 from C_2 $\frac{\pi}{\arccos \frac{r_1}{r_2}} \to \infty$. Since each polytope $P \in \mathcal{S}_f$ can contribute at most one line segment to $\mathcal{D}_f$, the size of the model necessary to represent a decision boundary that separates C_1 from C_2 also increases as the circles get closer together.

Now consider C_1^ϵ and C_2^ϵ under the L_2 norm, defined as $C_i^\epsilon = \{x \in \mathbb{R}^2 : \|x - C_i\|_2 \leq \epsilon\}$. Suppose that a fully connected network f described as above has sufficiently many parameters to represent a decision boundary that separates C_1 from C_2. Is f also capable of learning a *robust* decision boundary that separates C_1^ϵ from C_2^ϵ?

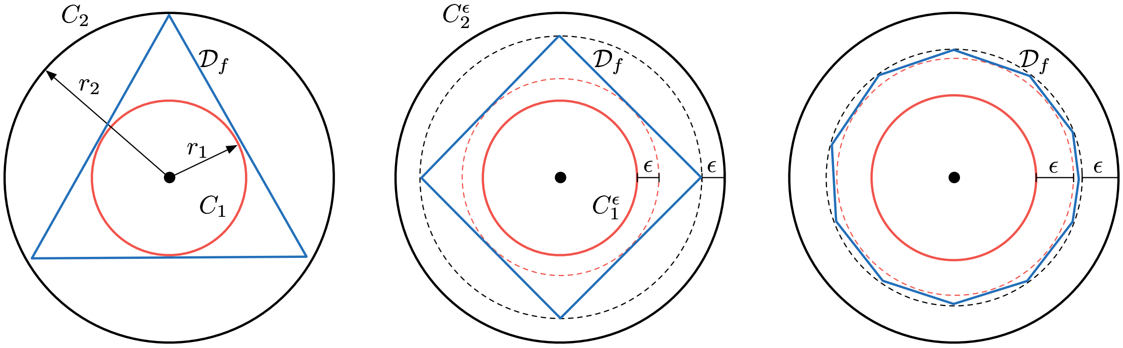

Figure 4.7: Separating two classes of data sampled from C_1 and C_2 may require a decision boundary $\mathcal{D}_f$ with only a few linear segments. However a decision boundary $\mathcal{D}_f$ that is robust to ϵ-perturbations must lie in gap between C_1^ϵ and C_2^ϵ. Learning a robust decision boundary may require more linear segments and thus a more expressive model. As we increase ϵ, demanding a more robust decision boundary, the gap between C_1^ϵ and C_2^ϵ decreases, and so the number of linear segments increases towards ∞.

For $\mathcal{D}_f$ to separate C_1^ϵ from C_2^ϵ it must lie in the region between C_1^ϵ and C_2^ϵ. In this setting each segment can contribute at most $2 \arccos \frac{r_1+\epsilon}{r_2-\epsilon}$ to the full 2π rotation around the origin. The minimum number of line segments that comprise a robust decision boundary $\mathcal{D}_f$ is lower bounded by $\frac{\pi}{\arccos \frac{r_1+\epsilon}{r_2-\epsilon}}$. As $\epsilon \to \frac{r_2-r_1}{2}$ this quantity approaches ∞. Even if f is capable of separating C_1 from C_2 we can choose ϵ such that $\frac{\pi}{\arccos \frac{r_1+\epsilon}{r_2+\epsilon}} \in \omega(h^{2\ell})$.

This simple example shows that learning decision boundaries that are robust to ϵ-adversarial examples may require substantially more powerful models than what is required to learn the original distributions. Furthermore the amount of additional resources necessary is dependent upon the amount of robustness required.

An Exponential Lower Bound

We present an exponential lower bound on the number of linear regions necessary to represent a decision boundary that is robust to L_2-perturbations of at most $\epsilon \leq \operatorname{rch}_2 \Lambda_2 - \tau$, in the simple case of two concentric $(d-1)$-spheres.

Theorem 30. *Let $S_1, S_2 \subset \mathbb{R}^d$ be two concentric $(d-1)$-spheres with radii $r_1 < r_2$ respectively and let $S = S_1 \cup S_2$. Let $f : \mathbb{R}^d \to \mathbb{R}$ be a fully connected neural network with ReLU activations. Suppose that f correctly classifies $S^{\operatorname{rch}_2 \Lambda_2 - \tau}$ for some $\tau \in [0, \operatorname{rch}_2 \Lambda_2]$. Said differently, the decision boundary of f lies in a τ-tubular neighborhood of the decision axis, $\mathcal{D}_f \subset \Lambda_2^\tau$. Then the number of linear regions N into which f subdivides $\mathbb{R}^d$ is lower bounded as*

$$N \geq 2\sqrt{\pi}\frac{\Gamma(\frac{d+1}{2})}{\Gamma(\frac{d}{2})}\left(\frac{r_1 + \operatorname{rch}_2 \Lambda_2}{4\tau}\right)^{\frac{d-1}{2}}. \tag{4.10}$$

Written asymptotically, $N \in \Omega\left(\frac{\sqrt{d}}{2^d}\left(\frac{r_1 + \operatorname{rch}_2 \Lambda_2}{\tau}\right)^{\frac{d-1}{2}}\right)$

Proof. For f to be robust to ϵ-adversarial examples for $\epsilon \leq \operatorname{rch}_2 \Lambda_2 - \tau$ the decision boundary $\mathcal{D}_f \subset \Lambda^\tau$. The boundary of Λ^τ is comprised of two disjoint $(d-1)$-spheres, which we will denote as $\partial\Lambda_1^\tau$ and $\partial\Lambda_2^\tau$ with radii $r_1 + \operatorname{rch}_2 \Lambda_2 - \tau$ and $r_1 + \operatorname{rch}_2 \Lambda_2 + \tau$ respectively. (It is standard in topology to use the ∂ symbol to denote the boundary of a topological space.)

The isoperimetric inequality states that a $(d-1)$-sphere minimizes the $(d-1)$-dimensional volume (thought of as "surface area") across all sets with fixed d-dimensional volume (thought of as "volume"). Since $\mathcal{D}_f \subset \Lambda^\tau$, the d-dimensional volume enclosed by $\mathcal{D}_f$ is at least as large as that of $\partial\Lambda_1^\tau$ and so we have that $\operatorname{surf} \partial\Lambda_1^\tau \leq \operatorname{surf} \mathcal{D}_f$.

Now consider any $(d-1)$-dimensional linear facet Π of the decision boundary $\mathcal{D}_f$. The normal space of Π is 1-dimensional; let $\boldsymbol{n}$ denote a unit vector orthogonal to Π. (There are two possible choices $\boldsymbol{n}$ and $-\boldsymbol{n}$.) Due to the spherical symmetry of Λ^τ and the fact that $\Pi \subset \Lambda^\tau$, the diameter of Π is maximized when Π is tangent to $\partial\Lambda_1^\tau$ at $(r_1 + \operatorname{rch}_2 \Lambda_2 - \tau)\boldsymbol{n}$ (or $-(r_1 + \operatorname{rch}_2 \Lambda_2 - \tau)\boldsymbol{n}$) and intersects $\partial\Lambda_2^\tau$. In pursuit of an upper bound, we will assume without loss of generality that Π has these properties. Let o denote the origin, $x = (r_1 + \operatorname{rch}_2 \Lambda_2 - \tau)\boldsymbol{n}$, and $y \in \Pi \cap \partial\Lambda_2^\tau$. We consider the right triangle $\triangle oxy$ with right angle at x. By basic properties of right triangles, $\frac{\operatorname{diam} \Pi}{2} \leq \|x - y\|_2 = \sqrt{(r_1 + \operatorname{rch}_2 \Lambda_2 + \tau)^2 - (r_1 + \operatorname{rch}_2 \Lambda_2 - \tau)^2} = \sqrt{4\tau(r_1 + \operatorname{rch}_2 \Lambda_2)}$. It follows that Π is contained in a $(d-1)$-dimensional ball of radius $\sqrt{4\tau(r_1 + \operatorname{rch}_2 \Lambda_2)}$. In particular the $(d-1)$-dimensional volume of Π is bounded as $\operatorname{vol}_{d-1}(\Pi) \leq \operatorname{vol}_{d-1} B(0, \sqrt{4\tau(r_1 + \operatorname{rch}_2 \Lambda_2)})$. The $(d-1)$-dimensional volume of $\mathcal{D}_f$ (again thought of as "surface area"), is equal to the sum of the $(d-1)$-dimensional volumes of the linear facets that comprise $\mathcal{D}_f$. Combining these inequalities gives the result.

$$\begin{aligned}
\frac{2\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2})}(r_1 + \mathrm{rch}_2\,\Lambda_2)^{d-1} = \mathrm{surf}\,\partial\Lambda_1^\tau &\le \mathrm{surf}\,\mathcal{D}_f \\
&\le N\,\mathrm{vol}_{d-1}\,B(0,\ \sqrt{4\tau(r_1 + \mathrm{rch}_2\,\Lambda_2)}) \\
&\le N\frac{\pi^{\frac{d-1}{2}}}{\Gamma(\frac{d+1}{2})}\,(4\tau(r_1 + \mathrm{rch}_2\,\Lambda_2))^{\frac{d-1}{2}} \\
2\sqrt{\pi}\frac{\Gamma(\frac{d+1}{2})}{\Gamma(\frac{d}{2})}\left(\frac{r_1 + \mathrm{rch}_2\,\Lambda_2}{4\tau}\right)^{\frac{d-1}{2}} &\le N
\end{aligned}$$

□

Prior work has experimentally verified that increasing the size of deep networks improves robustness [66]. Theorem 30 proves that there are settings in which robustness *requires* larger models.

4.8 Adversarial Examples from Voronoi Cells

Goodfellow et al. [41] originally proposed adversarial training where adversarial examples were constructed inside of an L_p-ball of radius ϵ. The use of the L_p-ball was meant to represent a simple notion of similarity between two images, delaying the complicated question of what is an adversarial image in favor of a tractable research problem. However it was never meant to be the final say on the threat model of the adversary and recent work has begun to explore alternative adversaries [53, 46].

In the previous sections we described a number of issues associated with the use of L_p-balls. Specifically we showed that the L_2-balls centered on a dense sample of $\mathcal{M}$ covers a negligible fraction of the neighborhood around $\mathcal{M}$. Thus, when constructing adversarial examples in the inner training loop, the adversary is restricted to constructing adversarial examples in a negligible fraction of the neighborhood around the data manifold. This vulnerability increases with the *codimension* $d - k$ of $\mathcal{M}$. Furthermore, for any p, a nearest neighbor classifier more effectively covers the neighborhood around $\mathcal{M}$ than a robust empirical risk minimization oracle, which outputs a classifier that is guaranteed to be correct in the L_p-balls centered on the data.

To remedy these shortcomings, we replace the L_p-ball constraint with a different geometric constraint, namely the Voronoi cell at each sample x, defined as

$$\mathrm{Vor}_p\,x = \{x' \in \mathbb{R}^d : \|x - x'\|_p \le \|z - x'\|_p\ \forall z \in X\backslash\{x\}\}. \tag{4.11}$$

In words, the Voronoi cell $\mathrm{Vor}_p\,x$ of x is the set of all points in $\mathbb{R}^d$ that are closer to x than to any other sample in X. The Voronoi diagram is defined as the collection of Voronoi cells, and their lower dimensional faces, for each sample in X. Figure 4.8 shows the Voronoi diagram for a dense sample from a dataset with two classes of data.

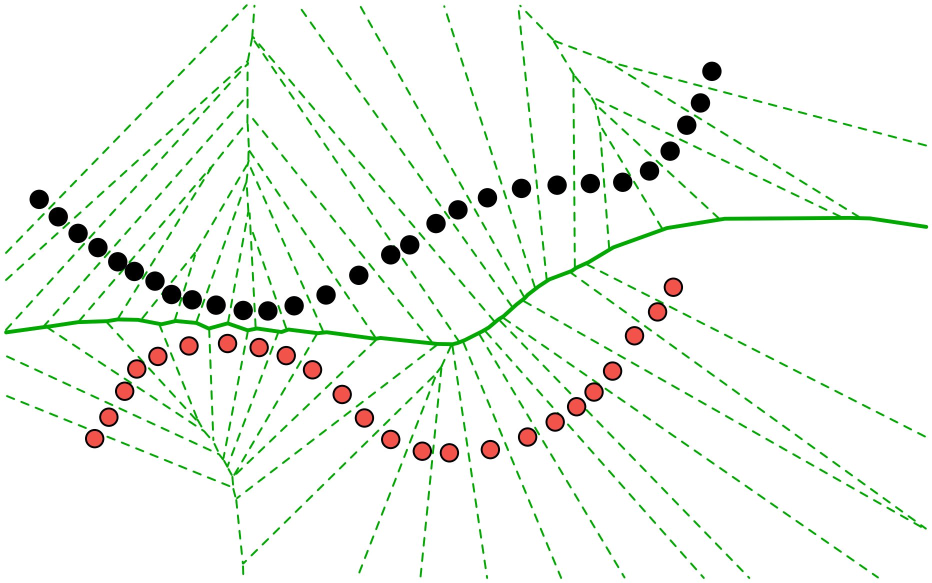

Figure 4.8: The Voronoi diagram for a dense sample drawn from a low-dimensional distribution with two classes, one in red and one in black. The Voronoi cells, shown in green, vary in size depending on how close a sample is to samples in the other class. The Voronoi edges that are adjacent to two samples from two different classes are shown in solid green, and approach a decision boundary which is as far from the data distribution as possible.

The Voronoi cell constraint has many advantages over the L_p-ball constraint. First the Voronoi cells *partition* the entirety of $\mathbb{R}^d$ and so the interiors of Voronoi cells generated by samples from different classes do not intersect. This is in contrast to L_p-balls which may intersect for sufficiently large ϵ. In particular the Voronoi cells partition the neighborhood around $\mathcal{M}$ and, for dense samples, are elongated in the directions normal to the data manifold [28]. Thus the Voronoi cells are well suited for high codimension settings. Second, the size of the Voronoi cells adapts to the data distribution. A Voronoi cell generated by a sample which is close to samples from a different class manifold is smaller, while those further away are larger. See Figure 4.8. Thus we do *not* need to set a value for ϵ in the optimization procedure. The constraint naturally adapts to the largest value of ϵ possible locally on the data manifold. Note that the maximum perturbation size possible will often vary as we move along the data manifold, and cannot be captured by a single number which, by necessity, is upper bounded by the smallest distance to a different class. In summary, the Voronoi constraint gives the adversary the freedom to explore the entirety of the neighborhood around $\mathcal{M}$.

At each iteration of standard adversarial training, we must solve the inner optimization problem $\max_{\delta \in B(0,\epsilon)} L(x+\delta, y; \theta)$ to generate an adversarial example. Goodfellow et al. [41] solve this problem using the fast gradient sign method (FGSM), while Madry et al. [66] use projected gradient descent. To incorporate Voronoi constraints, at each iteration of the outer training loop we must solve the inner optimization problem

$$\begin{aligned} \underset{\hat{x}}{\text{maximize}} \quad & L(\hat{x}, y; \theta) \\ \text{subject to} \quad & \|x - \hat{x}\|_p - \|z - \hat{x}\|_p \leq 0 \ \forall z \in X - \{x\}. \end{aligned} \tag{4.12}$$

When $p = 2$ the Voronoi cells are convex and so we can project a point onto a Voronoi cell by solving a quadratic program. Thus we can solve Problem 4.12 using projected gradient descent, as in [66]. When $p \neq 2$ the Voronoi cells are not necessarily convex. In this setting there are many approaches, such as barrier and penalty methods, one might employ to approximately solve Problem 4.12 [14].

However we found that the following heuristic is both fast and works well in practice. At each iteration of the outer training loop, for each training sample x in a batch, we generate adversarial examples by taking iterative steps in the direction of the gradient starting from x. Instead of projecting onto a constraint after each iterative step, we instead check if any of the Voronoi constraints of x shown in Equation 4.11 are violated. If no constraint is violated we perform the iterative update, otherwise we simply stop performing updates for x. Figure 4.9 illustrates the procedure.

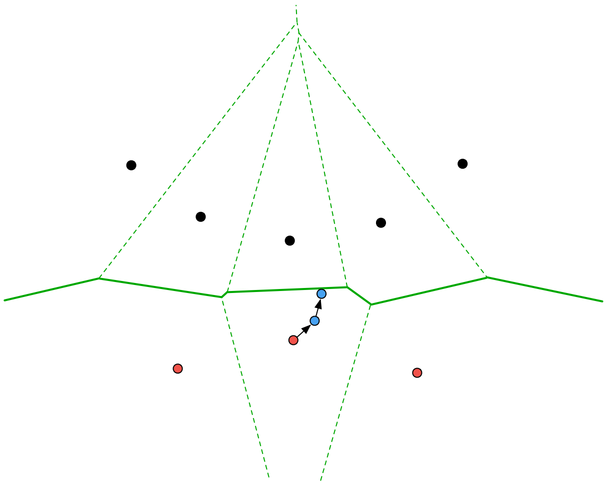

Figure 4.9: To construct an adversarial example within a Voronoi cell, we repeatedly take steps in the direction of the gradient of the loss, shown in blue. After each iteration we check if any of the Voronoi constraints are violated. We take the last iteration before a constraint is violated as our adversarial example.

Problem 4.12 has $n - 1$ constraints, one for each sample in $X \backslash \{x\}$. In practice however very few samples contribute to the Voronoi cell of x. Even fewer contribute to the faces of the Voronoi cell that are shared by samples in different classes, as shown in Figure 4.8. At each iteration, we perform a nearest neighbor search query to find the m nearest samples to x in each other class. That is we search for $m(C-1)$ samples where C is the number of classes. We do not impose constraints from samples in the same class as x; there is no benefit to restricting the adversary's movement with the neighborhood around the class manifold of x. In our experiments we set $m = 10$.

4.9 Adversarial Training with Voronoi Constraints

Madry et al. [66] formalize adversarial training by introducing the robust objective

$$\min_{\theta} \mathbb{E}_{(x,y)\in\mathcal{D}} \left[\max_{\hat{x}\in B(x,\epsilon)} L(\hat{x}, y; \theta) \right] \tag{4.13}$$

where $\mathcal{D}$ is the data distribution and B is a L_p-ball centered at x with radius ϵ. Their main contribution was the use of a strong adversary which used projected gradient descent to solve the inner optimization problem.

To incorporate Voronoi constraints, we replace the L_p-ball constraint in Equation 4.13 with the Voronoi cell at x. That is, we formalize the adversarial training objective as

$$\min_{\theta} \mathbb{E}_{(x,y)\in\mathcal{D}} \left[\max_{\hat{x}\in \mathrm{Vor}_p\, x} L(\hat{x}, y; \theta) \right], \tag{4.14}$$

where we use the optimization procedure described in Section 4.8 to solve the inner optimization problem.

4.10 Experiments

Datasets Section 4.6 suggests that as the codimension increases it should become easier to find adversarial examples. To verify this, we introduce two synthetic datasets, Circles and Planes, which allow us to carefully vary the codimension while maintaining dense samples. The Circles dataset consists of two concentric circles in the x_1-x_2-plane, the first with radius $r_1 = 1$ and the second with radius $r_2 = 3$, so that $\mathrm{rch}_2\, \Lambda_2 = 1$. We densely sample 1000 random points on each circle for both the training and the test sets. The Planes dataset consists of two 2-dimensional planes, the first in the $x_d = 0$ and the second in $x_d = 2$, so that $\mathrm{rch}_2\, \Lambda_2 = 1$. The first two axis of both planes are bounded as $-10 \leq x_1, x_2 \leq 10$, while $x_3 = \ldots = x_{d-1} = 0$. We sample the training set at the vertices of the grid described in Section 4.6, and the test set at the centers of the grid cubes, the blue points in Figure 4.4. Both planes are sampled so that the 1-tubular neighborhood X^1 covers the underlying planes, where X is the training set. The codimension for both datasets is $d-2$. We also evaluate on MNIST and CIFAR10.

Models Our controlled experiments on synthetic data consider a fully connected network with 1 hidden layer, 100 hidden units, and ReLU activations. We set the learning rate for Adam [59] as $\alpha = 0.1$. Our experimental results are averaged over 20 retrainings. For a fair comparison to adversarial training, our experiments on MNIST and CIFAR10 use the same model architectures as in [66]. We train the MNIST model using Adam for 100 epochs and the CIFAR10 model using SGD for 250 epochs.

Attacks On MNIST we apply 300-step projected gradient descent (PGD), with step sizes $\{0.05, 0.07, 0.1, 0.15, 0.17, 0.2\}$. On CIFAR10 we apply 20-step PGD with step sizes $\{2.0, 3.0, 4.0\}$. For both datasets we also apply the fast gradient sign method (FGSM) [41] to uncover possible gradient masking as recommended in [7].

Accuracy measures We plot the robust classification accuracy as a function of ϵ, for each of our datasets. Since one of the primary advantages of Voronoi constraints is that we do not need to set ϵ, we need a measure of robustness that considers the total robustness of the model. Thus we report the *normalized area under the curve* (NAUC) defined as

$$\mathrm{NAUC}(\mathrm{acc}) = \frac{1}{\epsilon_{\max}} \int_0^{\epsilon_{\max}} \mathrm{acc}(\epsilon)\, d\epsilon, \tag{4.15}$$

where $\mathrm{acc} : [0, \epsilon_{\max}] \to [0, 1]$ measures the classification accuracy and $\epsilon_{\max}$ is the largest perturbation considered. Note that $\mathrm{NAUC} \in [0, 1]$ with higher values corresponding to more robust models.

Implementation Details Constructing adversarial examples within the Voronoi cells, as described in Section 4.8, requires a nearest neighbor search query to find the m nearest samples to x in each other class. When the dataset remains constant throughout the course of training, this search can be performed once before training begins and reused at each iteration. However when the dataset is augmented during training, as in the case of data augmentation on CIFAR10, the nearest neighbor search query must be computed at each iteration. Since this computation is performed on the CPU, we create 16 threads, each with a copy of a k-d tree, which constantly pull mini-batches of samples from a queue and perform nearest neighbor queries. With 16 threads running in parallel, the bottleneck for training became the construction of adversarial examples on the GPU, and so adversarial training with Voronoi constraints ran in time similar to standard adversarial training.

High Codimension Reduces Robustness

Figure 4.10 (Left) shows the robustness of naturally trained networks on the Circles dataset as we increase the codimension. We see a steady decrease in robustness as we increase the codimension, on average.

Madry et al. [66] propose training against a projected gradient descent (PGD) adversary to improve robustness. Section 4.6 suggests that this should be insufficient to guarantee robustness, as X^ϵ is often a poor model for $\mathcal{M}^\epsilon$. We follow the adversarial training procedure of [66] by training against a PGD adversary with $\epsilon = 1$ under L_2-perturbations on the Planes dataset. Figure 4.10 (Right) shows that it is still easy to find adversarial examples for $\epsilon < 1$ and that as the codimension increases we can find adversarial examples for decreasing values of ϵ.

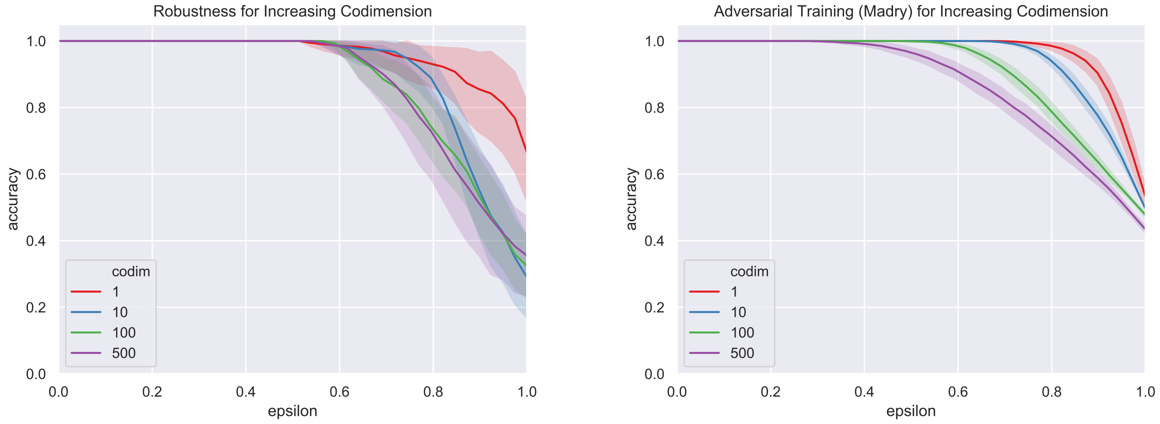

Figure 4.10: Left: As the codimension increases the robustness of decision boundaries learned by naturally trained networks on Circles decreases steadily. Right: Training using the adversarial training procedure of [66] is no guarantee of robustness; as the codimension increases it becomes easier to find adversarial examples for Planes.

In contrast, a nearest neighbor classifier achieves perfect robustness for all ϵ on both Circles and Planes. Nearest neighbors is robust even when the codimension is high, as long as the low-dimensional data manifold is well sampled. This is a consequence of the fact that the Voronoi cells of the samples are elongated in the directions normal to the data manifold when the sample is dense [28].

The Planes dataset is sampled so that the training set is a 1-cover of the underlying planes, which requires 450 sample points. Figure 4.11 shows the results of increasing the sampling density to a 0.5-cover (1682 samples) and a 0.25-cover (6498 samples). Increasing the sampling density improves the robustness of adversarial training at the same codimension and particularly in low-codimension. However adversarial training with a substantially larger training set does not produce a classifier as robust as a nearest neighbor classifier on a much smaller training set. Nearest neighbors is much more sample efficient than adversarial training, as predicted by Theorem 26.

Adversarial Perturbations are in the Directions Normal to the Data Manifold

Let η_x be an adversarial perturbation generated by FGSM with $\epsilon = 1$ at $x \in \mathcal{M}$. Note that the adversarial example is constructed as $\hat{x} = x + \eta_x$. In Figure 4.12 we plot a histogram of the angles $\angle(\eta_x, N_x\mathcal{M})$ between η_x and the normal space $N_x\mathcal{M}$ for the Circles dataset in codimensions $1, 10, 100,$ and 500. In codimension 1, 88% of adversarial perturbations make an angle of less than 10° with the normal space. Similarly in codimension 10, 97%, in codimension 100, 96%, and in codimension 500, 93%. As Figure 4.12 shows, nearly all adversarial perturbations make an angle less than 20° with the normal space. Our results are averaged over 20 retrainings of the model using

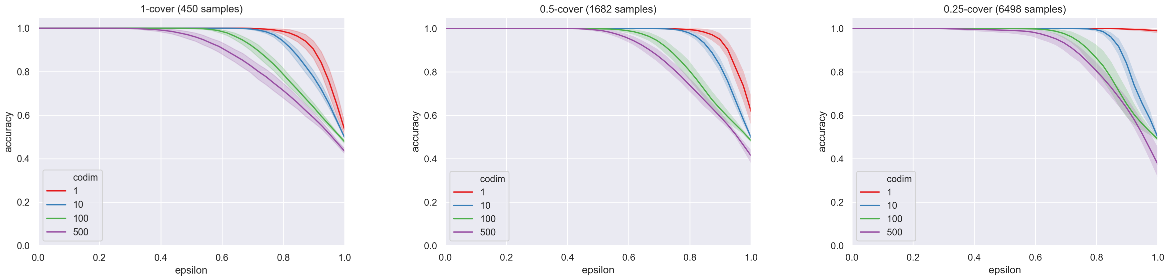

Figure 4.11: Adversarial training of [66] on the Planes dataset with a 1-cover (left), consisting of 450 samples, a 0.5-cover (center), 1682 samples, and a 0.25-cover (right), 6498 samples. Increasing the sampling density improves robustness at the same codimension. However even training on a significantly denser training set does not produce a classifier as robust as a nearest neighbor classifier on a much sparser training set.

SGD.

Throughout this chapter we've argued that high codimension is a key source of the pervasiveness of adversarial examples. Figure 4.12 shows that, when the underlying data manifold is well sampled, adversarial perturbations are well aligned with the normal space. When the codimension is high, there are many directions normal to the manifold and thus many directions in which to construct adversarial perturbations.

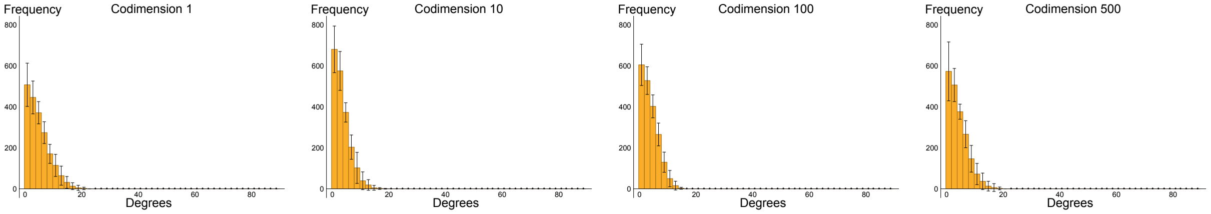

Figure 4.12: Histograms of the angle deviations of FGSM perturbations from the normal space for the Circles dataset in, from right to left, codimensions 1, 10, 100, and 500. Nearly all perturbations make an angle of less than 20° with the normal space.

Adversarial Training in High Codimensions

We showed that as the codimension of the Planes dataset increases, the adversarial training approach of Madry et al. [66] with training $\epsilon = 1$ became less robust. We believe that this is because the L_2-balls with radius 1 around the dataset covered an increasingly smaller fraction of the neighborhood around the data manifold.

Figure 4.13 shows that replacing the L_2 ball constraint with the Voronoi cells improves robustness in high codimension settings, on average. In codimension 10 (Figure 4.10 (Left)), our approach achieves NAUC of 0.99, while Madry's approach achieves NAUC of 0.94. In codimension 500 (Figure 4.13 (Right)), our approach achieves NAUC of 0.92, while Madry's approach achieves NAUC of 0.87.

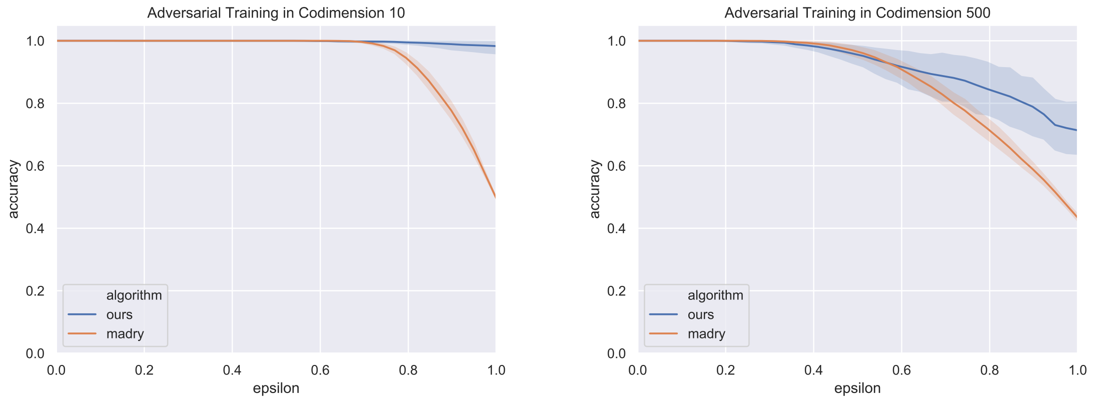

Figure 4.13: Adversarial training Voronoi constraints offers improved robustness in high codimension (10, 500) over standard adversarial training, on average.

MNIST and CIFAR10

To explore the performances of adversarial training with Voronoi constraints on more realistic datasets, we evaluate on MNIST and CIFAR10 and compare against the robust pretrained models of [66].[1,2]. We include the recently proposed Jacobian regularization algorithm of [47] with $\lambda_{\text{jr}} = 1.0$ as an additional baseline.

Figure 4.14 (Left) shows that our model maintains near identical robustness to the Madry model on MNIST up to $\epsilon = 0.3$, after which our model *significantly* outperforms the Madry model. The Madry model was explicitly trained for $\epsilon = 0.3$ perturbations. We emphasize that one advantage of our approach is that we did not need to set a value for the maximum perturbation size ϵ. The Voronoi cells adapt to the maximum size allowable locally on the data distribution. Our model maintains 76.3% accuracy at $\epsilon = 0.4$ compared to 2.6% accuracy for the Madry model. Furthermore our model achieves NAUC of 0.81, while the Madry model achieves NAUC of 0.67, an improvement of 20.8% and over the baseline. To our knowledge, this is the most robust MNIST model to L_∞ attacks.

[1] https://github.com/MadryLab/mnist_challenge
[2] https://github.com/MadryLab/cifar10_challenge

Figure 4.14 (Right) shows the results of our approach on CIFAR10. Both our model and the Madry model achieve NAUC of 0.29. However our approach trades natural accuracy for increased robustness against larger perturbations. This tradeoff is well-known and explored in [90, 48].

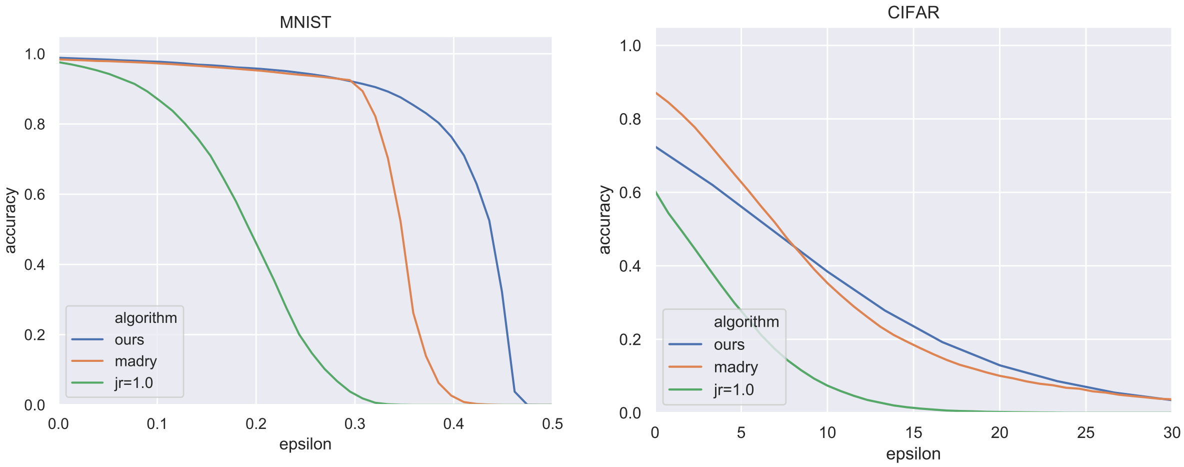

Figure 4.14: **Left:** Adversarial training with Voronoi constraints on MNIST. Our model has NAUC 0.81 and high classification accuracy after $\epsilon = 0.3$. In particular, our model maintains 76.3% accuracy at $\epsilon = 0.4$, compared to 2.6% accuracy for the Madry model. **Right:** On CIFAR10, both models achieve NAUC of 0.29, but our model trades natural accuracy for robustness to larger perturbations.

Increasing the Radius of the Norm Ball Constraint

A natural approach to improving the robustness of models produced by the adversarial training paradigm of [66] is to simply increase the maximum allowable perturbation size ϵ of the norm ball constraint. As shown in Figure 4.15, increasing the size of ϵ to 0.4, from the 0.3 with which [66] originally trained, and training for only 100 epochs produces a model which exhibits significantly worse robustness in the range $[0, 0.3]$ than the pretrained model. If we increase the number of training epochs to 150, the approach of [66] with $\epsilon = 0.4$ produces a model with improved robustness in the range $[0.3, 0.4]$, but that still exhibits the sharp drop in accuracy after 0.4. Additionally the model trained with $\epsilon = 0.4$ for 150 epochs performs worse than both the pretrained model and our model in the range $[0, 0.3]$. Our model achieves NAUC 0.81, while the model trained with $\epsilon = 0.4$ for 150 epochs achieves NAUC 0.76. We emphasize that our approach does not require us to set ϵ, which is particularly important in practice where the maximum amount of robustness achievable may not be known a-priori.

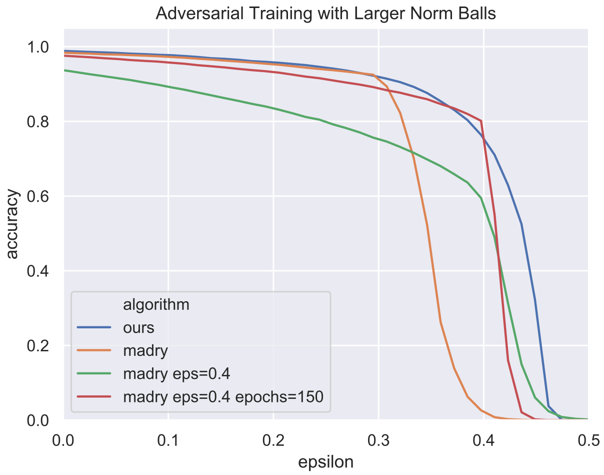

Figure 4.15: The adversarial training of [66] with $\epsilon = 0.4$ (shown in green) produces a model with significantly reduced robustness in the range $[0, 0.3]$. Increasing the number of epochs to 150, the resulting model (shown in red) does exhibit improved robustness in the range $[0.3, 0.4]$, at the expense of some robustness in the range $[0, 0.3]$ and still exhibits a sharp drop in accuracy after 0.4. The purple model achieves NAUC of 0.76, while our model achieves NAUC 0.81.

Chapter 5

Adaptive versus Standard Descent Methods and Robustness Against Adversarial Examples

5.1 Introduction

Adversarial examples are a pervasive phenomenon of machine learning models where perturbations of the input that are imperceptible to humans reliably lead to confident incorrect classifications [87, 41]. Since this phenomenon was first observed, researchers have attempted to develop methods which produce models that are robust to adversarial perturbations under specific attack models [94, 85, 72, 67, 66, 97]. As machine learning proliferates into society, including security-critical settings like health care [38] or autonomous vehicles [24], it is crucial to develop methods that allow us to understand the vulnerability of our models and design appropriate counter-measures.

Additionally there is a growing literature on the theory of adversarial examples. Many of these results attempt to understand adversarial examples by constructing examples of learning problems for which it is difficult to construct a classifier that is robust to adversarial perturbations. This difficulty may arise due to sample complexity [76], computational constraints [16, 27], or the high-dimensional geometry of the initial feature space [80, 56]. We expand upon these results in Section 5.2.

Currently less well-understood, and to our knowledge not addressed by the theoretical literature on adversarial examples, is how our algorithmic choices effect the robustness of our models. With respect to optimization and generalization, but importantly not robustness, the success of standard (or *non-adaptive*) gradient descent methods, including stochastic gradient descent (SGD) and SGD with momentum, is starting to be better understood [34, 1, 43, 42]. However, as an increasing amount of time has been spent training deep networks, researchers and practitioners have heavily employed *adaptive* gradient methods, such as Adam [59], Adagrad [35], and RMSprop [88], due to their rapid training times [54]. Unfortunately the properties of adaptive optimization algorithms are less well-understood than those of their non-adaptive counterparts. Wilson et al. [93] provide theoretical and empirical evidence which suggests that adaptive algorithms often produce solutions that generalize worse than those found by non-adaptive algorithms.

In this chapter, we study the robustness of solutions found by adaptive and non-adaptive algorithms to adversarial examples. Furthermore we study the effect of adversarial training on the geometry of the loss landscape and, consequently, on the solutions found by adaptive and non-adaptive algorithms for the adversarial training objective. This chapter makes the following contributions.

- We show an example of a learning problem for which the solution found by adaptive optimization algorithms exhibits qualitatively worse robustness properties against *both* L_2- and L_∞-adversaries than the solution found by non-adaptive algorithms. Furthermore the robustness of the adaptive solution decreases rapidly as the dimension of the problem increases, while the robustness of the non-adaptive solution is stable as the dimension increases.

- We fully characterize the geometry of the loss landscape of L_2-adversarial training in least-squares linear regression. The L_2-adversarial training objective $\mathcal{L}_2$ is convex everywhere; moreover, it is strictly convex everywhere except along either 0, 1, or 2 line segments, depending on the value of ϵ. Furthermore for nearly all choices[1] of ϵ, these line segments along which $\mathcal{L}_2$ is convex, but not strictly convex, lie outside of the rowspace and the gradient along these line segments is nonzero. It follows that any reasonable optimization algorithm finds the unique global minimum of $\mathcal{L}_2$.

- We conduct an extensive empirical evaluation to explore the effect of different optimization algorithms on robustness. Our experimental results suggest that non-adaptive methods consistently produce more robust models than adaptive methods.

- We provide a dataset consisting of 190 pretrained models on MNIST and CIFAR10 with various hyperparameter settings. Of these 190 pretrained models, 150 were used to find the best hyperparameter settings for our experiments and evaluated on a validation set. The remaining 40 pretrained models were evaluated on the test set. Of the 150 validation models, 88 were trained using natural training and 62 were trained using adversarial training. Of the 40 test models, 20 were trained using natural training and 20 were trained using adversarial training. They can be downloaded at https://www.dropbox.com/s/edfcnb97lzxl19z/models.zip.

5.2 Related Work

There has been a long line of work on the theory of adversarial examples. Schmidt et al. [76] explore the sample complexity required to produce robust models. They demonstrate a simple setting, a mixture of two Gaussians, in which a linear classifier with near perfect natural accuracy can be learned from a single sample, but *any* algorithm that produces *any* binary classifier requires $\Omega(\sqrt{d})$ samples to produce a robust classifier. Followup work by Bubeck et al. [16] suggests that adversarial

[1] For all $\epsilon \neq 1/\|X^\dagger y\|_2$

examples may arise from computational constraints. They exhibit pairs of distributions that differ only in a k-dimensional subspace, and are otherwise standard Gaussians, and show that while it is information-theoretically possible to distinguish these distributions, it requires exponentially many queries in the statistical query model of computation. We note that both of these constructions produce distributions whose support is the entirety of $\mathbb{R}^d$.

Bubeck et al. [16] further characterize five mutually exclusive "worlds" of robustness, inspired by similar characterizations in complexity theory [49]. A learning problem must fall into one of the following possibilities:

> **World 1**: No robust classifier exists, regardless of computational considerations or sample efficiency.
>
> **World 2**: Robust classifiers exists, but they are computationally inefficient to evaluate.
>
> **World 3**: Computationally efficient robust classifiers exist, but learning them requires more samples.
>
> **World 4**: Computationally efficient robust classifiers exist and can be learned from few samples, but learning is inefficient.
>
> **World 5**: Computationally efficient robust classifiers exists and can be learned efficiently from few samples.

While learning problems can be constructed that fall into each possible world, the question for researchers is into which world are problems from practice most likely to fall? Every theoretical construction, such as those by [76] and [16], can be thought of as providing evidence for the prevalence of one of the worlds. In the language of Bubeck et al. [16], the sampling complexity result of Schmidt et al. [76] provides evidence for World 3, by constructing an example of a problem that falls into world three. The learning problem constructed by Bubeck et al. [16] provides evidence for World 4. Subsequent work by Degwekar et al. [27] provides evidence for Worlds 2 and 4. Under standard cryptographic assumptions, Degwekar et al. [27] construct an a learning problem for which a computationally efficient non-robust classifier exists, no efficient robust classifier exists, but an inefficient robust classifier exists. Similarly, assuming the existence of one-way functions, they construct a learning problem for which an efficient robust classifier exists, but it is computationally inefficient to learn a robust classifier. Finally, in an attempt to understand how likely World 4 is in practice, they show that any task where an efficient robust classifier exists but is hard to learn in polynomial time implies one-way functions.[2]

Additionally there is a line of work that attempts to explain the pervasiveness of adversarial examples through the lens of high-dimensional geometry. Gilmer et al. [40] experimentally evaluated the setting of two concentric under-sampled 499-spheres embedded in $\mathbb{R}^{500}$, and concluded that adversarial examples occur on the data manifold. Shafahi et al. [80] suggest that adversarial examples may be an unavoidable consequence of the high-dimensional geometry of data. Their

[2] Thus at least one community will be happy.

result depends upon the use of an isopermetric inequality. The main drawback of these results, as well as the constructions of Schmidt et al. [76] and Bubeck et al. [16], is that they assume that the support of the data distribution has full or nearly full dimension. We do not believe this to be the case in practice. Instead we believe that the data distribution is often supported on a very low-dimensional subset of $\mathbb{R}^d$. This case is addressed in Khoury et al. [56], who consider the problem of constructing decision boundaries robust to adversarial examples when data is drawn from a low-dimensional manifold embedded in $\mathbb{R}^d$. They highlight the role of co-dimension, the difference between the dimension of the embedding space and the dimension of the data manifold, as a key source of the pervasiveness of adversarial vulnerability. Said differently, it is the low-dimensional structure of features embedded in high-dimensional space that contributes, at least in part, to adversarial examples. This idea is also explored by Nar et al. [68], but with emphasis on the cross-entropy loss.

We believe that problems in practice are most likely to fall into World 5, the best of all worlds. Problems in this class have robust classifiers which are efficient to evaluate and can be learned efficiently from relatively few samples. We simply haven't found the right algorithm for learning such classifiers. The goal of this chapter is to explore the effect of our algorithms on robustness. Specifically we wish to understand the robustness properties of solutions found by common optimization algorithms. To our knowledge no other work has explored the robustness properties of solutions found by different optimization algorithms.

5.3 Adaptive Algorithms May Significantly Reduce Robustness

Wilson et al. [93] explore the effect of different optimization methods on generalization both in a simple theoretical setting and empirically. For their main theoretical result, they construct a learning problem for which the solution found by *any* adaptive method, denoted w_{ada}, has worse generalization properties than the solution found by non-adaptive methods, denoted w_{SGD}. We recall their construction in the next subsection. We describe the adaptive solution w_{ada} and the non-adaptive solution w_{SGD}.

Generalization and robustness are different properties of a classifier. A classifier can generalize well but have terrible robustness properties, as we often see in practice. On the other hand, a constant classifier generalizes poorly, but has perfect robustness [97]. Wilson. et al. [93] study the *generalization* properties of w_{ada} and w_{SGD}, but not their robustness properties. In the fourth subsection we study the robustness properties of w_{ada} and w_{SGD}. Specifically, we show that w_{SGD} exhibits superior robustness properties to w_{ada} against *both* L_2- and L_∞-adversaries.

A Simple Learning Problem

Let $X \in \mathbb{R}^{n \times d}$ be a design matrix representing a dataset with n sample points and d features and let $y \in \{\pm 1\}^n$ be a vector of labels. Wilson et al. [93] restrict their attention to binary classification problems of this type, and learn a classifier by minimizing the least-squares loss

$$\min_{w} \mathcal{L}(X, y; w) = \min_{w} \frac{1}{2}\|Xw - y\|_2^2. \tag{5.1}$$

They construct the following learning problem for which they can solve for both the adaptive and non-adaptive solutions in closed form. Their construction uses an infinite-dimensional feature space for simplicity, but they note that $6n$ dimensions suffice. For $i \in 1 \ldots n$, sample $y_i = 1$ with probability p, and $y_i = -1$ with probability $1 - p$ for some $p > 0.5$. Then set x_i to be the infinite-dimensional vector

$$x_{ij} = \begin{cases} y_i & j = 1 \\ 1 & j = 2, 3 \\ 1 & j = 4 + 5(i-1) \\ (1 - y_i)/2 & j = 5 + 5(i-1), \ldots, 8 + 5(i-1) \\ 0 & \text{otherwise.} \end{cases} \tag{5.2}$$

For example, a dataset with three sample points following Equation 5.3 is

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \ldots \\ -1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & \ldots \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & \ldots \end{pmatrix}. \tag{5.3}$$

The first feature encodes the label, and is alone sufficient for classification. Note that this trick of encoding the label is also commonly used in the robustness literature to construct examples of hard-to-learn-robustly problems [16, 27]. The second and third feature are identically 1 for every sample. Then there is a subset of five dimensions which are identified with x_i and contain a set of features which are *unique* to x_i. If $y_i = 1$ then there is a single 1 in this subset of five dimensions and x_i is the only sample with a 1 in any of these five positions. If $y_i = -1$ then all five dimensions are set to 1 and again x_i is the only sample with a 1 at these five positions.

While this problem may seem contrived, it contains several properties that are common in machine learning problems and that are particularly important for robustness. It contains a single robust feature that is strongly correlated with the label. However it may not be easy for an optimization algorithm to identify such a feature. Additionally there are many non-robust features which are weakly or not at all correlated with the label, but which may appear useful for generalization because they are uniquely identified with samples from specific classes. Wilson et al. [93] show that both adaptive and non-adaptive methods find classifiers that place at least some weight on every nonzero feature.

The Adaptive Solution w_{ada}

Let (X, y) be generated by the generative model in Section 5.3. When initialized at the origin, Wilson et al. [93] show that *any* adaptive optimization algorithm – such as RMSprop, Adam, and Adagrad – minimizing Equation 5.1 for (X, y) converges to $w_{\text{ada}} \propto v$ where

$$v_j = \begin{cases} 1 & j = 1 \\ 1 & j = 2, 3 \\ y_{\lfloor (j+1)/5 \rfloor} & j > 3 \text{ and } x_{\lfloor (j+1)/5 \rfloor, j} = 1 \\ 0 & \text{otherwise.} \end{cases} \tag{5.4}$$

Thus we can write $w_{\text{ada}} = \tau v$ for some positive constant $\tau > 0$. On a test example $(x_{\text{test}}, y_{\text{test}})$, that is distinct from all the training examples, $\langle w_{\text{ada}}, x_{\text{test}} \rangle = \tau(y_{\text{test}} + 2) > 0$. Thus w_{ada} labels every unseen example as a positive example.

The Non-adaptive Solution w_{SGD}

For (X, y), let $\mathcal{P}, \mathcal{N}$ denote the sets of positive and negative samples in X respectively. Let $n_+ = |\mathcal{P}|, n_- = |\mathcal{N}|$ and note that $n = n_+ + n_-$. When the weight vector is initialized in the row space of X, Wilson et al. [93] show that all non-adaptive methods – such as gradient descent, SGD, SGD with momentum, Nesterov's method, and conjugate gradient – converge to $w_{\text{SGD}} = X^\dagger y$, where $X^\dagger$ denotes the pseudo-inverse of X. That is, among the infinitely many solutions of the underdetermined system $Xw = y$, non-adaptive methods converge to the solution which minimizes $||w||_2$, and thus maximizes the L_2-margin. Specifically $w_{\text{SGD}} = \sum_{i \in \mathcal{P}} \alpha_+ x_i + \sum_{j \in \mathcal{N}} \alpha_- x_j$ where

$$\alpha_+ = \frac{4n_- + 5}{15n_+ + 3n_- + 8n_+n_- + 5}, \quad \alpha_- = -\frac{4n_+ + 1}{15n_+ + 3n_- + 8n_+n_- + 5}.$$

Note that these values for α_+, α_- differ slightly from those presented in [93]. In Appendix C we discuss in detail two errors in their derivation that lead to this discrepancy. These errors do not qualitatively change their results. Furthermore, in the proof of Theorem 33 we carefully discuss under what conditions $\langle w_{\text{SGD}}, x_{\text{test}} \rangle$ is positive and negative for $y_{\text{test}} = \pm 1$. For now, we simply state that for all $n_+, n_- \geq 1$, w_{SGD} correctly classifies every test example.

Analyzing the Robustness of w_{ada} and w_{SGD}

In this section we analyze the robustness properties of w_{ada} and w_{SGD} against L_2- and L_∞-adversaries. We show that w_{SGD} exhibits considerably more robustness against *both* L_2- and L_∞-adversaries than w_{ada}. A priori this is surprising; one may have expected w_{ada}, which is a small L_∞-norm solution, to be more robust to L_∞-perturbations, while w_{SGD}, which is a small L_2-norm solution, to be robust to L_2-perturbations. However this expectation is wrong. Interestingly the robustness of w_{ada} against both L_2- and L_∞-adversaries decreases as the dimension increases, whereas the robustness of w_{SGD} does not. Finally, neither method recovers the "obvious" solution $w^* = (1, 0, \ldots, 0)$, which generalizes well and is optimally robust against both L_2- and L_∞-perturbations.

Theorems 31 and 33 are our main results of this section. We start by computing the robustness of w_{ada} against L_2- and L_∞-adversaries.

Theorem 31. *Let $(x_{\text{test}}, y_{\text{test}})$ be a test sample that is correctly classified by w_{ada} and let $\delta \in \mathbb{R}^d$ be a perturbation. The adaptive solution w_{ada} is robust against any L_2-perturbation for which*

$$||\delta||_2 < \frac{\sqrt{9n_+ + 1125n_- + 27}}{25n_- + n_+ + 3} \tag{5.5}$$

and any L_∞-perturbation for which

$$||\delta||_\infty < \frac{3}{3 + n_+ + 5n_-}. \tag{5.6}$$

Furthermore these bounds are tight, meaning that a closed L_2- or L_∞-ball with these radii centered at x_{test} intersects the decision boundary.

Proof. Let δ be an adversarial perturbation and let x_{test} be a test sample. Then

$$\begin{aligned}
\langle w_{\text{ada}}, x_{\text{test}} + \delta\rangle &= \langle w_{\text{ada}}, x_{\text{test}}\rangle + \langle w_{\text{ada}}, \delta\rangle \\
&= \tau(y_{\text{test}} + 2) + \langle w_{\text{ada}}, \delta\rangle \\
&= \tau(y_{\text{test}} + 2) + \tau\left(\delta_1 + \delta_2 + \delta_3 + \sum_{i\in\mathcal{P}} \delta_i - 5\sum_{j\in\mathcal{N}} \delta_j\right) \\
&= 3\tau + \tau\left(\delta_1 + \delta_2 + \delta_3 + \sum_{i\in\mathcal{P}} \delta_i - 5\sum_{j\in\mathcal{N}} \delta_j\right) \\
&= \tau(y_{\text{test}} + 2) - \tau\delta\,(3 + n_+ + 5n_-)
\end{aligned}$$

The second last equality follows from the fact that x_{test} is correctly classified by w_{ada}, and so $y_{\text{test}} = 1$. Notice that to flip the sign of the classifier using the smallest L_∞-perturbation, it is optimal to distribute the magnitude of the perturbation equally to each δ_i, where the signs of each δ_i are -1 for $i \in \{1, 2, 3\} \cup \mathcal{P}$ and $+1$ for $i \in \mathcal{N}$. It follows that to flip the sign of the classifier requires

$$\begin{aligned}
\delta\,(3 + n_+ + 5n_-) &> 3 \\
\delta &> \frac{3}{3 + n_+ + 5n_-}.
\end{aligned}$$

To find the smallest L_2-perturbation we must instead solve the constrained optimization problem

$$\begin{aligned}
\min_{\delta} \quad & \sum_i \delta_i^2 \\
\text{s.t.} \quad & \left(\delta_1 + \delta_2 + \delta_3 + \sum_{i\in\mathcal{P}} \delta_i - 5\sum_{j\in\mathcal{N}} \delta_j\right) < -3
\end{aligned} \tag{5.7}$$

where $R^2 = \sum_i \delta_i^2$ is the squared-radius of the smallest L_2-ball that crosses the decision boundary. The Lagrangian for this problem is

$$\mathcal{L}(\delta, \lambda) = \sum_i \delta_i^2 + \lambda\left(\delta_1 + \delta_2 + \delta_3 + \sum_{i\in\mathcal{P}} \delta_i - 5\sum_{j\in\mathcal{N}} \delta_j + 3\right).$$

The partial derivatives are

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \delta_i} &= \begin{cases} 2\delta_i + \lambda & i = 1, 2, 3 \text{ or } i \in \mathcal{P} \\ 2\delta_i - 5\lambda & i \in \mathcal{N} \end{cases} \\
\frac{\partial \mathcal{L}}{\partial \lambda} &= \delta_1 + \delta_2 + \delta_3 + \sum_{i \in \mathcal{P}} \delta_i - 5 \sum_{j \in \mathcal{N}} \delta_j + 3.
\end{aligned}$$

Setting the first set of partial derivatives to 0 gives

$$\delta_i = \begin{cases} -\frac{\lambda}{2} & i = 1, 2, 3 \text{ or } i \in \mathcal{P} \\ \frac{5\lambda}{2} & i \in \mathcal{N} \end{cases}, \tag{5.8}$$

which can then be used to solve the last equation $\frac{\partial \mathcal{L}}{\partial \lambda} = 0$ yielding

$$\lambda = \frac{6}{25n_- + n_+ + 3}.$$

Substituting the expression for λ back into Equation 5.8 gives

$$\delta_i = \begin{cases} \frac{-3}{25n_- + n_+ + 3} & i = 1, 2, 3 \text{ or } i \in \mathcal{P} \\ \frac{15}{25n_- + n_+ + 3} & i \in \mathcal{N} \end{cases}. \tag{5.9}$$

Then the minimum L_2-perturbation R is derived from

$$\begin{aligned}
R^2 &= \sum_i \delta_i^2 \\
&= (3 + n_+)\left(\frac{-3}{25n_- + n_+ + 3}\right)^2 + 5n_-\left(\frac{15}{25n_- + n_+ + 3}\right)^2 \\
&= \frac{9(n_+ + 3) + 1125n_-}{(25n_- + n_+ + 3)^2} \\
R &= \frac{\sqrt{9n_+ + 1125n_- + 27}}{25n_- + n_+ + 3}.
\end{aligned}$$

□

Corollary 32. *Asymptotically, the L_2- and L_∞-robustness of w_{ada} are, respectively,*

$$\Theta\left(\frac{1}{\sqrt{n_+ + n_-}}\right) \textit{ and } \Theta\left(\frac{1}{n_+ + n_-}\right).$$

In particular both the L_2- and L_∞-robustness go to 0 as the number of samples $n_+, n_- \to \infty$.

Corollary 32 makes clear, qualitatively, the result in Theorem 31. The rate at which the L_2- and L_∞-robustness of w_{ada} decrease reflects a dependence on dimension. The number of dimensions on which w_{ada} puts nonzero weight increases as we increase the number of samples, which reduces robustness. We also find it interesting that, despite classifying every test point as a positive example, w_{ada}'s predictions on correctly classified test samples are brittle. In summary, w_{ada} exhibits nearly no robustness against L_2- or L_∞-adversaries.

Next we show that w_{SGD} exhibits significant robustness against *both* L_2- and L_∞-adversaries.

Theorem 33. *Let $(x_{\text{test}}, y_{\text{test}})$ be a test sample that is correctly classified by w_{SGD} and let $\delta \in \mathbb{R}^d$ be a perturbation. The SGD solution w_{SGD} is robust against any L_2-perturbation for which*

$$||\delta||_2 \leq \begin{cases} \frac{15n_+ + 8n_+n_- - n_-}{\sqrt{64n_+^2n_-^2 + 160n_+^2n_- + 75n_+^2 + 32n_+n_-^2 + 60n_+n_- + 70n_+ + 3n_-^2 + 5n_-}} & y_{\text{test}} = 1 \\ \frac{-5n_+ + 8n_+n_- + 3n_-}{\sqrt{64n_+^2n_-^2 + 160n_+^2n_- + 75n_+^2 + 32n_+n_-^2 + 60n_+n_- + 70n_+ + 3n_-^2 + 5n_-}} & y_{\text{test}} = -1 \end{cases} \tag{5.10}$$

and any L_∞-perturbation for which

$$||\delta||_\infty \leq \begin{cases} \frac{15n_+ + 8n_+n_- - n_-}{20n_+ + 32n_+n_- + 4n_-} & y_{\text{test}} = 1 \\ \frac{-5n_+ + 8n_+n_- + 3n_-}{20n_+ + 32n_+n_- + 4n_-} & y_{\text{test}} = -1. \end{cases} \tag{5.11}$$

Furthermore these bounds are tight, meaning that a closed L_2- or L_∞-ball with these radii centered at x_{test} intersects the decision boundary.

Proof. It is worth taking a moment to understand $\langle w_{\text{SGD}}, x_{\text{test}}\rangle$ when $y_{\text{test}} = 1$ and when $y_{\text{test}} = -1$. In particular, it will be important in our proofs to understand the signs of *each* term.

First, we have $\alpha_+ > 0$ and $\alpha_- < 0$ by definition. When $y_{\text{test}} = 1$ we have

$$\begin{aligned} \langle w_{\text{SGD}}, x_{\text{test}}\rangle &= (n_+\alpha_+ - n_-\alpha_-) + 2(n_+\alpha_+ + n_-\alpha_-) \\ &= \frac{5n_+ + n_- + 8n_+n_-}{15n_+ + 3n_- + 8n_+n_- + 5} + \frac{2(5n_+ - n_-)}{15n_+ + 3n_- + 8n_+n_- + 5} \\ &= \frac{15n_+ + 8n_+n_- - n_-}{15n_+ + 3n_- + 8n_+n_- + 5}. \end{aligned}$$

The denominator is clearly positive, so w_{SGD} correctly classifies x_{test} so long as $15n_+ + 8n_+n_- - n_- > 0$, which is true for any $n_+, n_- \geq 1$.

When $y_{\text{test}} = -1$ we have

$$\begin{aligned} \langle w_{\text{SGD}}, x_{\text{test}}\rangle &= -(n_+\alpha_+ - n_-\alpha_-) + 2(n_+\alpha_+ + n_-\alpha_-) \\ &= -\frac{5n_+ + n_- + 8n_+n_-}{15n_+ + 3n_- + 8n_+n_- + 5} + \frac{2(5n_+ - n_-)}{15n_+ + 3n_- + 8n_+n_- + 5} \\ &= \frac{5n_+ - 8n_+n_- - 3n_-}{15n_+ + 3n_- + 8n_+n_- + 5}. \end{aligned}$$

In this case, w_{SGD} correctly classifies x_{test} so long as $5n_+ - 8n_+n_- - 3n_- < 0$, which is true for any $n_+, n_- \geq 1$. Thus w_{SGD} correctly classifies every test example so long as there at least one training example from each class.

We will also be interested in the signs of the individual terms in $\langle w_{\text{SGD}}, x_{\text{test}} \rangle$. Note that $5n_+ + n_- + 8n_+n_- > 0$ for any $n_+, n_- \geq 1$, and so $(n_+\alpha_+ - n_-\alpha_-)$ is positive. Lastly $5n_+ - n_- > 0$ so long as $n_+ > n_-/5$, and so $(n_+\alpha_+ + n_-\alpha_-) > 0$ if and only if $n_+ > n_-/5$. For convenience we will assume that $n_+ > n_-/5$ from here onward which will allow us to consider fewer cases.

Let δ be an adversarial perturbation and let x_{test} be a test sample. Then

$$\langle w_{\text{SGD}}, x_{\text{test}} + \delta \rangle = \langle w_{\text{SGD}}, x_{\text{test}} \rangle + \langle w_{\text{SGD}}, \delta \rangle$$

where

$$\langle w_{\text{SGD}}, x_{\text{test}} \rangle = y_{\text{test}}(n_+\alpha_+ - n_-\alpha_-) + 2(n_+\alpha_+ + n_-\alpha_-)$$

and

$$\langle w_{\text{SGD}}, \delta \rangle = (n_+\alpha_+ - n_-\alpha_-)\delta_1 + (n_+\alpha_+ + n_-\alpha_-)(\delta_2 + \delta_3) + \alpha_+ \sum_{i \in \mathcal{P}} \delta_i + \alpha_- \sum_{j \in \mathcal{N}} \left(\delta_{j,1} + \ldots + \delta_{j,5}\right).$$

There are two cases to consider corresponding to $y_{\text{test}} = \pm 1$.

Suppose that $y_{\text{test}} = 1$. To flip the sign we need $\langle w_{\text{SGD}}, \delta \rangle < -\langle w_{\text{SGD}}, x_{\text{test}} \rangle$. For brevity's sake, we define

$$C \equiv (n_+\alpha_+ - n_-\alpha_-)\delta_1 + (n_+\alpha_+ + n_-\alpha_-)(\delta_2 + \delta_3) + \alpha_+ \sum_{i \in \mathcal{P}} \delta_i + \alpha_- \sum_{j \in \mathcal{N}} \left(\delta_{j,1} + \ldots + \delta_{j,5}\right) + \langle w_{\text{SGD}}, x_{\text{test}} \rangle.$$

The constraint $C < 0$ is equivalent to $\langle w_{\text{SGD}}, \delta \rangle < -\langle w_{\text{SGD}}, x_{\text{test}} \rangle$. We can ensure the sign of $\langle w_{\text{SGD}}, \delta \rangle$ is negative by choosing each δ_i opposite in sign to the term by which it is multiplied in C. Note that, by our assumptions on n_+, n_-, $(n_+\alpha_+ - n_-\alpha_-), (n_+\alpha_+ + n_-\alpha_-), \alpha_+ > 0$ and $\alpha_- < 0$. Thus we choose $\text{sign}(\delta_{j,1\ldots,5}) = 1$ for all $j \in \mathcal{N}$ and $\text{sign}(\delta_i) = -1$ otherwise.

For a perturbation in the L_∞-norm, the optimal solution sets each δ_i to the same magnitude, and so to change the sign the perturbation δ must be at least

$$\begin{aligned}
\delta &> \frac{\langle w_{\text{SGD}}, x_{\text{test}} \rangle}{(n_+\alpha_+ - n_-\alpha_-) + 2(n_+\alpha_+ + n_-\alpha_-) + n_+\alpha_+ - 5n_-\alpha_-} \\
&= \frac{\langle w_{\text{SGD}}, x_{\text{test}} \rangle}{4(n_+\alpha_+ - n_-\alpha_-)} \\
&= \frac{3n_+\alpha_+ + n_-\alpha_-}{4(n_+\alpha_+ - n_-\alpha_-)} \\
&= \frac{15n_+ + 8n_+n_- - n_-}{20n_+ + 32n_+n_- + 4n_-}.
\end{aligned}$$

Now suppose that $y_{\text{test}} = -1$. In this case to flip the sign we need $\langle w_{\text{SGD}}, \delta \rangle > -\langle w_{\text{SGD}}, x_{\text{test}} \rangle$, (equivalently $C > 0$). Note that in this case $\langle w_{\text{SGD}}, x_{\text{test}} \rangle$ is negative, and so we choose the signs of each δ_i to match the signs of the terms by which δ_i is multiplied. We choose $\text{sign}(\delta_{j,1\ldots,5}) = -1$ and $\text{sign}(\delta_i) = 1$ otherwise. Thus to change the sign the perturbation δ must satisfy

$$\begin{aligned}\delta &> \frac{-\langle w_{\text{SGD}}, x_{\text{test}}\rangle}{(n_+\alpha_+ - n_-\alpha_-) + 2(n_+\alpha_+ + n_-\alpha_-) + n_+\alpha_+ - 5n_-\alpha_-} \\ &= \frac{-\langle w_{\text{SGD}}, x_{\text{test}}\rangle}{4(n_+\alpha_+ - n_-\alpha_-)} \\ &= \frac{-n_+\alpha_+ - 3n_-\alpha_-}{4(n_+\alpha_+ - n_-\alpha_-)} \\ &= \frac{-5n_+ + 8n_+n_- + 3n_-}{20n_+ + 32n_+n_- + 4n_-}.\end{aligned}$$

To find the smallest L_2-perturbation, in the case where $y_{\text{test}} = 1$, we must solve the constrained optimization problem

$$\begin{aligned}\min_{\delta} \quad & \sum_i \delta_i^2 \\ \text{s.t.} \quad & C \leq 0\end{aligned} \tag{5.12}$$

where $R^2 \equiv \sum_i \delta_i^2$ is the squared-radius of the smallest L_2-ball that touches the decision boundary. The Lagrangian for this problem is

$$\mathcal{L}(\delta, \lambda) = \sum_i \delta_i^2 + \lambda C.$$

The partial derivatives are

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial \delta_i} &= \begin{cases} 2\delta_i + \lambda(n_+\alpha_+ - n_-\alpha_-) & i = 1 \\ 2\delta_i + \lambda(n_+\alpha_+ + n_-\alpha_-) & i = 2, 3 \\ 2\delta_i + \lambda\alpha_+ & i \in \mathcal{P} \\ 2\delta_{i,j} + \lambda\alpha_- & i \in \mathcal{N}, j \in [5] \end{cases} \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= C.\end{aligned}$$

Setting the first set of partial derivatives to 0 gives

$$\delta_i = \begin{cases} -\frac{\lambda}{2}(n_+\alpha_+ - n_-\alpha_-) & i = 1 \\ -\frac{\lambda}{2}(n_+\alpha_+ + n_-\alpha_-) & i = 2, 3 \\ -\frac{\lambda}{2}\alpha_+ & i \in \mathcal{P} \\ -\frac{\lambda}{2}\alpha_- & i \in \mathcal{N}, j \in [5] \end{cases}, \tag{5.13}$$

which can then be used to solve the last equation $\frac{\partial \mathcal{L}}{\partial \lambda} = C = 0$ yielding

$$\lambda = \frac{\langle w_{\text{SGD}}, x_{\text{test}}\rangle}{\frac{1}{2}(n_+\alpha_+ - n_-\alpha_-)^2 + (n_+\alpha_+ + n_-\alpha_-)^2 + \frac{1}{2}n_+\alpha_+^2 + \frac{5}{2}n_-\alpha_-^2}.$$

Substituting the expression for λ back into Equation 5.13 and solving for R gives

$$\begin{aligned}
R^2 &= \sum_i \delta_i^2 \\
&= \frac{\lambda^2}{4}\left((n_+\alpha_+ - n_-\alpha_-)^2 + 2(n_+\alpha_+ + n_-\alpha_-)^2 + n_+\alpha_+^2 + 5n_-\alpha_-^2\right) \\
&= \frac{\lambda^2}{2}\left(\frac{1}{2}(n_+\alpha_+ - n_-\alpha_-)^2 + (n_+\alpha_+ + n_-\alpha_-)^2 + \frac{1}{2}n_+\alpha_+^2 + \frac{5}{2}n_-\alpha_-^2\right) \\
&= \frac{\langle w_{\text{SGD}}, x_{\text{test}}\rangle^2}{(n_+\alpha_+ - n_-\alpha_-)^2 + 2(n_+\alpha_+ + n_-\alpha_-)^2 + n_+\alpha_+^2 + 5n_-\alpha_-^2} \\
R &= \frac{\langle w_{\text{SGD}}, x_{\text{test}}\rangle}{\sqrt{(n_+\alpha_+ - n_-\alpha_-)^2 + 2(n_+\alpha_+ + n_-\alpha_-)^2 + n_+\alpha_+^2 + 5n_-\alpha_-^2}} \\
&= \frac{3n_+\alpha_+ + n_-\alpha_-}{\sqrt{(n_+\alpha_+ - n_-\alpha_-)^2 + 2(n_+\alpha_+ + n_-\alpha_-)^2 + n_+\alpha_+^2 + 5n_-\alpha_-^2}} \\
&= \frac{15n_+ + 8n_+n_- - n_-}{\sqrt{64n_+^2n_-^2 + 160n_+^2n_- + 75n_+^2 + 32n_+n_-^2 + 60n_+n_- + 70n_+ + 3n_-^2 + 5n_-}}
\end{aligned}$$

The case with $y_{\text{test}} = -1$ is similar, but with the constraint $-C \leq 0$, which yields a similar solution for λ, except that the numerator is $-\langle w_{\text{SGD}}, x_{\text{test}}\rangle > 0$. Subsequently

$$R = \frac{-5n_+ + 8n_+n_- + 3n_-}{\sqrt{64n_+^2n_-^2 + 160n_+^2n_- + 75n_+^2 + 32n_+n_-^2 + 60n_+n_- + 70n_+ + 3n_-^2 + 5n_-}}.$$

□

Corollary 34. *Asymptotically, the L_2- and L_∞-robustness of w_{SGD} are both $\Theta(1)$. In particular the L_2-robustness approaches 1 and the L_∞-robustness approaches $\frac{1}{4}$ as the number of samples $n_+, n_- \to \infty$.*

Unsurprisingly, w_{SGD}, which maximizes the L_2-margin, exhibits near-optimal robustness against L_2-adversaries. As the number of samples increases, the L_2-robustness of w_{SGD} approaches 1. Perhaps surprisingly, w_{SGD} also exhibits moderate robustness to L_∞-perturbations. As the number of samples increases, the L_∞-robustness of w_{SGD} approaches $\frac{1}{4}$. Unlike w_{ada}, the amount of robustness exhibited by w_{SGD} does not decrease as the dimension increases, instead asymptotically approaching a constant.

However the L_2-robustness of w_{SGD} is not exactly 1 for any finite sample. One class ($y_{\text{test}} = 1$) approaches 1 from above, while the other class ($y_{\text{test}} = -1$) approaches 1 from below. To maximize the margin, w_{SGD} places a small amount of weight on every other nonzero feature, even though all but the first are useless for classification. This lack of sparsity is also what causes the L_∞-robustness to drop from a possible maximum of 1 to $\frac{1}{4}$. In contrast, $w^* = (1, 0, \ldots, 0)$ generalizes perfectly, has L_2-robustness equal to 1 for both classes, and, as an added benefit, has L_∞-robustness equal to 1 for both classes. Thus we have an example of a problem for which the max L_2-margin solution could reasonably be considered to not be the best classifier against L_2-perturbations.

Furthermore, w^* is *not* in the row space of X. (w_{SGD} is the projection of w^* onto the row space.) Thus non-adaptive methods, when restricted to the row space, are *incapable* of recovering w^*, irrespective of sample complexity [76] or computational considerations [16]. This is simply the wrong algorithm for the desired objective. In the next section we study the effect that adversarial training has on the loss landscape and on the solutions found by various optimization algorithms.

5.4 Adversarial Training (Almost) Always Helps

In the previous section we presented a learning problem for which adaptive optimization methods find a solution with significantly worse robustness properties against both L_2- and L_∞-adversaries compared to non-adaptive methods. In this section we consider a different algorithm, adversarial training, for finding robust solutions to Equation 5.1. We are interested in two questions. First, does adversarial training sufficiently regularize the loss landscape so that adaptive and non-adaptive methods find solutions with identical or qualitatively similar robustness properties? Second, are the solutions to the robust objective qualitatively different than those found by natural training or does adversarial training simply choose a robust solution from the space of solutions to the natural problem? We address the first question in the following two subsections for L_2-adversarial training and the second question in the third subsection for the learning problem defined in Section 5.3.

The Adversarial Training Objective

Madry et al. [66] formalize adversarial training by introducing the robust objective

$$\min_w \mathbb{E}_{(x,y)\in\mathcal{D}}\left[\max_{\delta\in\Delta} \mathcal{L}(x+\delta, y; w)\right] \tag{5.14}$$

where $\mathcal{D}$ is the data distribution, Δ is a perturbation set meant to enforce a desired constraint, and $\mathcal{L}$ is a loss function. The goal then is to find a setting of the parameters w of the model that minimize the expected loss against the worst-case perturbation in Δ.

Take $\mathcal{L}$ as in Equation 5.1 and Δ to be an L_p-ball of radius $\epsilon > 0$. In the linear case, we can solve the inner maximization problem exactly.

$$\begin{aligned}
\max_{\{\delta_i\}_{i\in[n]}\in\Delta^n} \mathcal{L}(x+\delta_i, y; w) &= \max_{\{\delta_i\}_{i\in[n]}\in\Delta^n} \frac{1}{2}\sum_{i=1}^n (\langle x_i+\delta_i, w\rangle - y_i)^2 && (5.15)\\
&= \max_{\{\delta_i\}_{i\in[n]}\in\Delta^n} \frac{1}{2}\sum_{i=1}^n \Big((\langle x_i, w\rangle - y_i)^2 + 2\langle\delta_i, w\rangle(\langle x_i, w\rangle - y_i) + \langle\delta_i, w\rangle^2\Big)\\
&= \frac{1}{2}\sum_{i=1}^n \Big((\langle x_i, w\rangle - y_i)^2 + 2\epsilon\|w\|_* \operatorname{sign}(\langle x_i, w\rangle - y_i)(\langle x_i, w\rangle - y_i) + \epsilon^2\|w\|_*^2\Big)\\
&= \frac{1}{2}\|Xw-y\|_2^2 + \epsilon\|w\|_*\|Xw-y\|_1 + \frac{\epsilon^2 n}{2}\|w\|_*^2. && (5.16)
\end{aligned}$$

The third identity follows from the definition of the dual norm, where $\|\cdot\|_*$ denotes the norm dual to the L_p norm that defines Δ. As a technical note, it is important that $\text{sign}(0) = 1$ (or -1) and *not* equal to 0. This choice represents the fact that the solution to the inner maximization problem for each individual squared term $(\langle x_i + \delta_i, w\rangle - y_i)^2$ is nonzero even if $x_i^\top w - y_i = 0$.

At first glance the objective looks similar to ridge regression or Lasso, particularly when we consider L_2- and L_∞-adversarial training for which the dual norms are L_2 and L_1 respectively. However the solutions to this objective are not, in general, identical to the ridge regression or Lasso solutions. In Section 5.4 we will show how the second term $\epsilon\|w\|_*\|Xw - y\|_1$ influences the geometry of the loss landscape when $\|\cdot\|_* = \|\cdot\|_2$.

The Geometry of the Loss Landscape

For the remainder of the chapter we will *exclusively* analyze the case where Δ is an L_2-ball, leaving the case of L_∞ for future work. We define the loss of interest

$$\mathcal{L}_2(X, y; w) = \frac{1}{2}\|Xw - y\|_2^2 + \epsilon\|w\|_2\|Xw - y\|_1 + \frac{\epsilon^2 n}{2}\|w\|_2^2. \tag{5.17}$$

To build intuition, suppose that $Xw = y$ is an underdetermined system. (Our results will not depend on this assumption.) The set of solutions is given by the affine subspace $S = \{X^\dagger y + u : u \in \text{nullspace}(X)\}$, where $X^\dagger$ is the pseudo-inverse. The first thing to notice about $\epsilon\|w\|_2\|Xw - y\|_1$ is that, on its own, it is non-convex, having local minima both at the origin and in S. Along any straight path starting at the origin and ending at a point in S, the loss landscape induced by $\epsilon\|w\|_2\|Xw - y\|_1$ is negatively curved.

The second thing to notice about $\epsilon\|w\|_2\|Xw - y\|_1$ is that it is non-smooth. To understand the loss landscape of Equation 5.17, it is crucial to understand where $\|Xw - y\|_1$ is non-smooth. The term $\|Xw - y\|_1 = \sum_i |x_i^\top w - y_i|$ is non-smooth at any point w where some $x_i^\top w - y_i = 0$. Geometrically, $x_i^\top w - y_i = 0$ is the equation of a hyperplane h_i with normal vector x_i and bias y_i. The hyperplane h_i partitions $\mathbb{R}^d$ into two halfspaces h_i^+, h_i^- such that every point $w \in h_i^+$ has $\text{sign}(x_i^\top w - y_i) = 1$ and $w \in h_i^-$ has $\text{sign}(x_i^\top w - y_i) = -1$. The set of hyperplanes $\{h_i : i \in [n]\}$ define a *hyperplane arrangement* $\mathcal{H}$, a subdivision of $\mathbb{R}^d$ into convex cells. See Figure 5.1. Let $C \in \mathcal{H}$ be a cell of the hyperplane arrangement. (We use "cell" to refer to a d-dimensional face $\mathcal{H}$. When considering a lower dimensional face of $\mathcal{H}$ we will refer to the dimension explicitly.) Every point w in the interior $\text{Int}\, C$ of C lies on the same side of every hyperplane h_i as every other point in $\text{Int}\, C$. Thus we can identify each C with a *signature* $s = \text{sign}(Xw - y)$ for any $w \in \text{Int}\, C$.

Theorem 35, our main result of this section, fully characterizes the geometry of Equation 5.17.

Theorem 35. *$\mathcal{L}_2$ is always a convex function, whose optimal solution(s) always lies in* rowspace(X). *There are four possible cases, three of which depend on the value of ϵ.*

1. *If $Xw = y$ is an inconsistent system, then $\mathcal{L}_2$ is a strictly convex function.*

2. *If $Xw = y$ is a consistent system and $\epsilon \in (0, 1/\|X^\dagger y\|_2)$ then $\mathcal{L}_2$ is a convex function. Moreover, $\mathcal{L}_2$ is strictly convex everywhere except along two line segments, both of which have one*

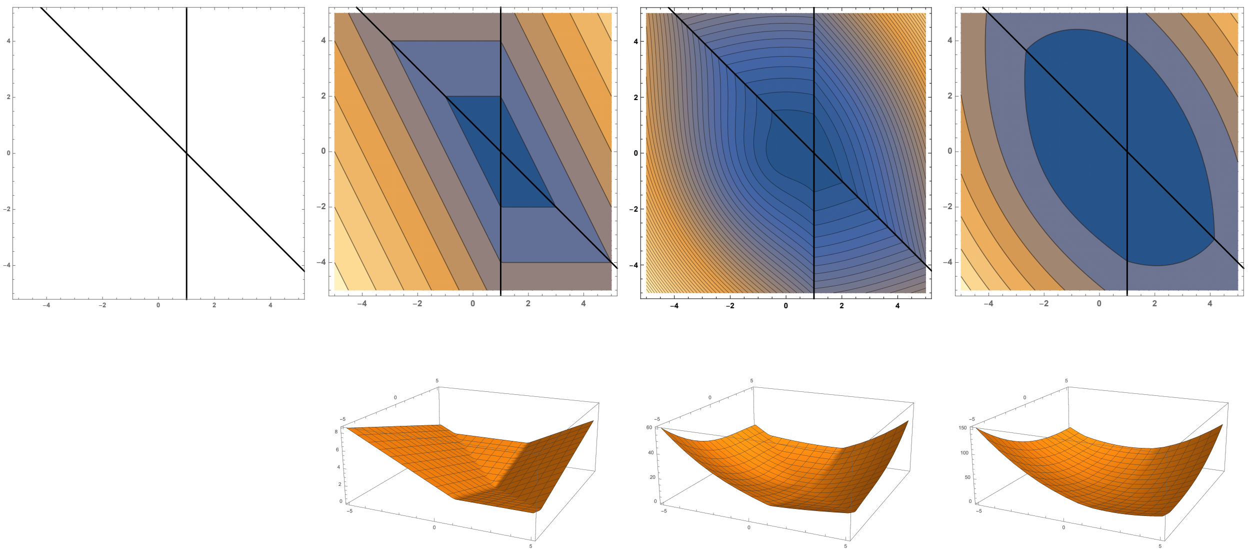

Figure 5.1: The top leftmost figure shows a hyperplane arrangement with two lines that subdivide $\mathbb{R}^2$ into four convex cells. The top center left figure shows the isocontours of $||Xw - y||_1$. Within each convex cell, the isocontours behave as the linear function $s^\top(Xw - y)$, where s is the signature of the cell. The isocontours are non-smooth along the two black lines. The top center right figure shows the isocontours of $\epsilon||w||_2||Xw - y||_1$, which are clearly non-convex. The top rightmost figure shows the isocontours of the function $\mathcal{L}_2 = \frac{1}{2}||Xw - y||_2^2 + \epsilon||w||_2||Xw - y||_1 + \frac{\epsilon^2 n}{2}||w||_2^2$. Notice how these isocontours are convex and non-smooth along the two black lines and the asymmetry of the isocontours caused by the $\epsilon||w||_2||Xw - y||_1$ term. The bottom row shows the graphs of the functions in the top row.

endpoint at the origin and terminate at $X^\dagger y \pm u$ respectively, for some nonzero $u \in$ nullspace(X). *Thus both line segments lie outside of* rowspace(X). *The gradient at every point on these line segments is nonzero, and so the optimal solution is unique and found in* rowspace(X) *at a point of strict convexity.*

3. *If $Xw = y$ is a consistent system and $\epsilon = 1/||X^\dagger y||_2$, then $\mathcal{L}_2$ is a convex function. Moreover, $\mathcal{L}_2$ is strictly convex everywhere except along a single line segment with one endpoint at the origin and the other endpoint at $X^\dagger y$. The optimal solution(s) may or may not lie along this line.*

4. *If $Xw = y$ is a consistent system and $\epsilon > 1/||X^\dagger y||_2$, then $\mathcal{L}_2$ is a strictly convex function.*

Furthermore, $\mathcal{L}_2$ is subdifferentiable everywhere.

Proof. Lemma 67 states that $\mathcal{L}_2$ is convex and that transitions between cells are strictly convex. The cases in the theorem statement correspond to the cases in Lemma 66 which describe the geometry of the cell containing the origin. Finally Lemma 65 states that any optimal solution must be in the rowspace of X. Lemma 68 states that $\mathcal{L}_2$ is subdifferentiable everywhere. □

The proof of Theorem 35 first characterizes the geometry of $\mathcal{L}_2$ restricted to the interior of each convex cell $C \in \mathcal{H}$. We denote this function by $\mathcal{L}_2|_{\text{Int}\,C}$. If the signature s of C is not equal to $-y$, meaning that C does not contain the origin, then $\mathcal{L}_2|_{\text{Int}\,C}$ is always strongly convex. The cell C with signature $s = -y$ is the only cell in which $\mathcal{L}_2$ might not be strongly convex. The cases in Theorem 35 correspond to the cases that characterize the geometry of $\mathcal{L}_2|_{\text{Int}\,C}$ for $s = -y$. The transitions between cells are strictly convex and the subdifferential is non-empty at these transitions.

We find it very interesting that $\mathcal{L}_2$ is strictly or strongly convex almost everywhere. For $\epsilon \in (0, 1/\|X^\dagger y\|_2)$, $\mathcal{L}_2$ is convex, but fails to be strictly convex only along two line segments which lie outside of the rowspace of X. The gradient along these line segments is nonzero, and so this particular type of convexity does not prevent an optimization algorithm from finding the *unique* solution, which is in the rowspace of X. For $\epsilon = 1/\|X^\dagger y\|_2$, $\mathcal{L}_2$ is convex, but not strictly convex, only along a single line segment in the rowspace of X. However the condition $\epsilon \neq 1/\|X^\dagger y\|_2$ can be ensured by an infinitesimal perturbation. The following remark is immediate.

Corollary 36. *Suppose $\epsilon \neq 1/\|X^\dagger y\|_2$. Then any optimization algorithm which is guaranteed to find or converge to the global minimum for a strictly convex subdifferentiable function and which does not prematurely terminate at a point with nonzero gradient finds the unique global minimum of $\mathcal{L}_2$.*

Corollary 36 states that, in the linear case, any reasonable optimization algorithm finds the unique global optimum of $\mathcal{L}_2$, almost always.[3] We conclude that, in the linear case, L_2-adversarial training does indeed sufficiently regularize the loss landscape so that any optimization algorithm finds the same solution.

Unfortunately Theorem 35 does not give a closed-form expression for the solution(s) of $\mathcal{L}_2$. In Lemma 68 we characterize the subdifferential $\partial\mathcal{L}_2(w)$ at every point. However solving for w where $0 \in \partial\mathcal{L}_2(w)$ is similar to solving a linear program, and so we suspect that no closed-form solution exists. In Section 5.4 we discuss the solution for $\mathcal{L}_2$ in the particular case of the learning problem defined in Section 5.3 and show that the max-margin solution is often *not* the solution recovered by adversarial training.

The Solutions to the Learning Problem of Section 5.3

While we know of no technique to characterize the set of solutions for $\mathcal{L}_2$ in general, we can still make some statements about the solution in specific instances, such as the learning problem

[3] The condition on "premature termination" in Corollary 36 is meant to rule out the following case. One could construct an optimization algorithm that is guaranteed to converge for strictly convex functions, but terminates early upon detecting a point at which there exists a direction in which the function is convex but not strictly convex, even if the gradient is nonzero and the global minimum is at a point of strict convexity. We doubt any commonly used optimization algorithm would have difficulty with the geometry of $\mathcal{L}_2$.

described in Section 5.3. First, since the minimizer(s) of $\mathcal{L}_2$ must lie in the rowspace, the "obvious" solution w^* to the learning problem in Section 5.3 is not recovered by L_2-adversarial training. A priori, one might guess that the minimum L_2-norm solution $X^\top\alpha$ is the solution to $\mathcal{L}_2$. However this is only true under specific conditions which depend on the class imbalance.

Theorem 37. *Let (X, y) be the learning problem defined in Section 5.3. $X^\top\alpha$ is a solution to $\mathcal{L}_2$ if and only if*

$$\epsilon \leq \frac{\sqrt{64n_+^2n_-^2 + 160n_+^2n_- + 75n_+^2 + 32n_+n_-^2 + 60n_+n_- + 70n_+ + 3n_-^2 + 5n_-}}{\max\left\{4n_-^2 + 4n_-n_+ + 5n_+ + 5n_-, 4n_+^2 + 4n_-n_+ + n_+ + n_-\right\}}. \tag{5.18}$$

Let $c > 0$ be a constant such that $n_+ = cn_-$. If $\epsilon \leq \min\left\{\frac{2c}{1+c}, \frac{2}{1+c}\right\}$ then $X^\top\alpha$ is a solution to $\mathcal{L}_2$.

Proof. The gradient at the minimum L_2-norm solution $X^\top\alpha$ is

$$\begin{aligned}\nabla\mathcal{L}_2(X^\top\alpha) &= X^\top(XX^\top\alpha - y) + \epsilon\|X^\top\alpha\|_2X^\top s + \epsilon\|XX^\top\alpha - y\|_1\frac{X^\top\alpha}{\|X^\top\alpha\|_2} + \epsilon^2nX^\top\alpha \\ &= \epsilon\|X^\top\alpha\|_2X^\top s + \epsilon^2nX^\top\alpha.\end{aligned}$$

Setting $\nabla\mathcal{L}_2 = 0$ gives

$$\begin{aligned}-X^\top s &= \epsilon n\frac{X^\top\alpha}{\|X^\top\alpha\|_2} \\ -\sum_{i=1}^n s_ix_i &= \frac{\epsilon n}{\|X^\top\alpha\|_2}\left(\sum_{i\in\mathcal{P}}\alpha_+x_i + \sum_{j\in\mathcal{N}}\alpha_-x_j\right).\end{aligned} \tag{5.19}$$

Lemma 68 states that, at $X^\top\alpha$, there exists a subgradient for every choice of $s \in [-1, 1]^n$. To prove the result we must show that there exists some choice of s that satisfies Equation 5.19, which we will do by showing that, under the condition on ϵ, the coefficient of each x_i on the right-hand side of Equation 5.19 is in the range $[-1, 1]$.

Since $s_i \in [-1, 1]$ the negative sign on the left-hand side of Equation 5.19 is inconsequential. It is sufficient to show that $\frac{\epsilon n\alpha_+}{\|X^\top\alpha\|_2}$ and $\frac{\epsilon n\alpha_-}{\|X^\top\alpha\|_2}$ are in the range $[-1, 1]$. Necessity follows from the fact that the rows of X are linearly independent.

$$\begin{aligned}\frac{\epsilon n\alpha_+}{\|X^\top\alpha\|_2} &= \epsilon\frac{(n_+ + n_-)\alpha_+}{\sqrt{(n_+\alpha_+ - n_-\alpha_-)^2 + 2(n_+\alpha_+ + n_-\alpha_-)^2 + n_+\alpha_+^2 + 5n_-\alpha_-^2}} \\ &= \epsilon\frac{4n_-^2 + 4n_-n_+ + 5n_+ + 5n_-}{\sqrt{64n_+^2n_-^2 + 160n_+^2n_- + 75n_+^2 + 32n_+n_-^2 + 60n_+n_- + 70n_+ + 3n_-^2 + 5n_-}} \\ &\leq 1\end{aligned}$$

where the last inequality follows from the condition on ϵ. Note also that $\frac{\epsilon n\alpha_+}{\|X^\top\alpha\|_2} \geq 0$ by definition. The case for α_- is similar.

Now assume that $n_+ = cn_-$. The right hand side of Equation 5.18 becomes

$$\frac{\sqrt{64c^2n_-^4 + 160c^2n_-^3 + 75c^2n_-^2 + 32cn_-^3 + 60cn_-^2 + 70cn_- + 3n_-^2 + 5n_-}}{\max\left\{4n_-^2 + 4cn_-^2 + 5cn_- + 5n_-, 4c^2n_-^2 + 4cn_-^2 + cn_- + n_-\right\}}. \tag{5.20}$$

The maximum evaluates as

$$\max\left\{4n_-^2 + 4cn_-^2 + 5cn_- + 5n_-, 4c^2n_-^2 + 4cn_-^2 + cn_- + n_-\right\} = \begin{cases} 4n_-^2 + 4cn_-^2 + 5cn_- + 5n_- & \text{if } c < \frac{1+n_-}{n_-} \\ 4c^2n_-^2 + 4cn_-^2 + cn_- + n_- & \text{if } c \geq \frac{1+n_-}{n_-}. \end{cases}$$

Within each of these ranges it can be checked, using Mathematica, that the gradient of Equation 5.20 is negative. Thus we can consider the limit as $n_- \to \infty$, which gives the lower bound

$$\frac{\sqrt{64c^2n_-^4 + 160c^2n_-^3 + 75c^2n_-^2 + 32cn_-^3 + 60cn_-^2 + 70cn_- + 3n_-^2 + 5n_-}}{\max\left\{4n_-^2 + 4cn_-^2 + 5cn_- + 5n_-, 4c^2n_-^2 + 4cn_-^2 + cn_- + n_-\right\}} \geq \min\left\{\frac{2c}{1+c}, \frac{2}{1+c}\right\}.$$

Taking $\epsilon \leq \min\left\{\frac{2c}{1+c}, \frac{2}{1+c}\right\}$ is a sufficient condition, but not necessary due to the gap in the lower bound. □

While we need the first condition, which is both necessary and sufficient, to draw our forthcoming conclusion, the second, merely sufficient, condition provides greater intuition. The first condition states that $X^\top\alpha$ is a solution if and only if ϵ is sufficiently small, as a function of n_+, n_-. For the learning problem in Section 5.3, we know that $\epsilon = 1$ is achievable, so it is natural to ask how large an ϵ is allowable by Equation 5.18. This relationship depends on the class imbalance, and so we set $n_+ = cn_-$ and derive the condition in the second part of the proof of Theorem 37, which is a lower bound on the right-hand side of Equation 5.18. The term $\min\left\{\frac{2c}{1+c}, \frac{2}{1+c}\right\}$ is at most 1 when $c = 1$, but can be arbitrarily less than 1 depending on c; see Figure 5.2. We note also that the gap between the lower bound $\min\left\{\frac{2c}{1+c}, \frac{2}{1+c}\right\}$ and the right-hand side of Equation 5.18 is already small for $n_- \approx 20$ and vanishes as $n_- \to \infty$. Thus we conclude that if ϵ is close to 1 and the dataset is even moderately imbalanced, $X^\top\alpha$, which maximizes the L_2-margin, is *not* a solution for $\mathcal{L}_2$.

5.5 Experiments

In this section we experimentally explore the effect of different optimization algorithms on robustness for deep networks. We are interested in the following questions. (1) Do different optimization algorithms give qualitatively different robustness results? (2) Does adversarial training reduce or eliminate the influence of the optimization algorithm? (3) Do adaptive or non-adaptive methods consistently outperform the others for both L_2- and L_∞-adversarial attacks, even if the difference is small?

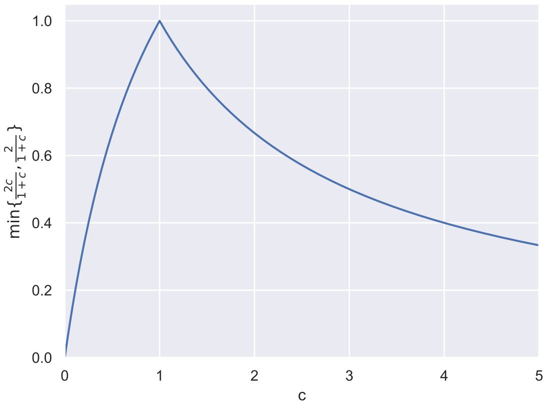

Figure 5.2: The bound $\min\left\{\frac{2c}{1+c}, \frac{2}{1+c}\right\}$ as a function of the class imbalance c. If the classes are perfectly balanced ($c = 1$), then we can take ϵ up to 1 and recover the minimum L_2-norm solution $X^\top\alpha$. As the class imbalance increases the maximum allowable ϵ for which we recover $X^\top\alpha$ decreases rapidly.

Models For MNIST our model consist of two convolutional layers with 32 and 64 filters respectively, each followed by 2×2 max pooling. After the two convolutional layers, there are two fully connected layers each with 1024 hidden units. For CIFAR10 we use a ResNet18 model [45]. We use the same model architectures for both natural and adversarial training. These models were chosen because they are small enough for us to run a large hyperparameter search.

Parameters for Adversarial Training For adversarial training we use the approach of Madry et al. [66], and train against a projected gradient descent (PGD) adversary. For MNIST with L_∞-adversarial training, we train against a 40-step PGD adversary with step size 0.01 and maximum perturbation size of $\epsilon = 0.3$. For L_2-adversarial training we train against a 40-step PGD adversary with step size 0.05 and maximum perturbation size of $\epsilon = 1.5$. For CIFAR10 with L_∞-adversarial training, we train against against a 10-step PGD adversary with step size 0.007 ($= 2/255$) and maximum perturbation size of $\epsilon = 0.031$ ($= 8/255$). For L_2-adversarial training we train against a 10-step PGD adversary with step size 0.039 ($= 10/255$) and a maximum perturbation size of $\epsilon = 0.117$ ($= 30/255$).

Attacks for Evaluation On MNIST we apply 100-step PGD with 10 random restarts. For L_∞ we apply PGD with step sizes $\{0.01, 0.05, 0.1, 0.2\}$, and for L_2 we apply PGD with step sizes $\{0.05, 0.1, 10^5, 10^9\}$. On CIFAR10 we apply 20-step PGD with 5 random restarts. For L_∞ we apply PGD with step sizes $\{2/255, 3/255, 4/255\}$ and for L_2 we apply PGD with step sizes $\{10/255, 20/255, 10^5\}$. We also apply the gradient-free BoundaryAttack++ [17]. We evaluate these attacks *per sample*, meaning that if any attack successfully constructs an adversarial example for a sample x at a specific ϵ, it reduces the robust accuracy of the model at that ϵ.

Metrics We plot the robust classification accuracy for each attack as a function of $\epsilon \in [0, \epsilon_{\text{max}}]$. We are interested in both natural and adversarial training. Usually when heuristic methods for adversarial training are evaluated, they are compared at the specific ϵ for which the model was adversarially trained. Such a comparison is arbitrary for naturally trained models and is also unsatisfying for adversarially trained models. To compare the robustness of different optimization algorithms we instead consider the area under the robustness curve. Following Khoury et al. [55], we report the *normalized area under the curve* (NAUC) defined as

$$\text{NAUC}(\text{acc}) = \frac{1}{\epsilon_{\text{max}}} \int_0^{\epsilon_{\text{max}}} \text{acc}(\epsilon)\, d\epsilon, \tag{5.21}$$

where $\text{acc} : [0, \epsilon_{\text{max}}] \to [0, 1]$ measures the classification accuracy. Note that $\text{NAUC} \in [0, 1]$, with higher values corresponding to more robust models.

Hyperparameter Selection We perform an extensive hyperparameter search over the learning rate and (if applicable) momentum parameter(s) of each optimization algorithm to identify parameter settings that produce the most robust models. For each dataset we set aside a validation set of size 5000 from the training set. We then train models, with the architecture described above, for each of the hyperparameter settings described below for 100 epochs. All optimization algorithms are started from the same initialization. We evaluate the robustness of each model as described above on the validation set. The hyperparameter settings for each optimization algorithm that achieve the largest NAUC are then used to train new models and then evaluated on the full test set. These final results are the ones we report in this section. The hyperparameters we explored were influenced by the hyperparameter search of Wilson et al. [93].

The following search space is defined for MNIST. For SGD we consider the learning rates $\{2, 1, 0.5, 0.1, 0.01, 0.001, 0.003, 0.0001\}$. For SGD with momentum we consider the set of learning rates for SGD for each of the momentum settings $\{0.9, 0.8, 0.7\}$. For Adam, Adagrad, and RMSprop we consider initial learning rates $\{0.1, 0.01, 0.001, 0.003, 0.0001\}$.

The following search space is defined for CIFAR10. For SGD we consider the learning rates $\{2, 1, 0.5, 0.25, 0.1\}$. For SGD with momentum we consider the set of learning rates for SGD for each of the momentum settings $\{0.9, 0.8, 0.7\}$. For both SGD and SGD with momentum we reduce the learning rate using the ReduceLROnPlateau scheduler in PyTorch. For Adam and RMSprop we consider the initial learning rates $\{0.005, 0.001, 0.0005, 0.0003, 0.0001, 0.00005\}$. For Adagrad we consider initial learning rates $\{0.1, 0.05, 0.01, 0.0075, 0.0005\}$.

Unfortunately adversarially training takes an order of magnitude longer than natural training, since in the innermost loop we must perform an iterative PGD attack to construct adversarial examples. Due to our limited resources, we consider only a subset of the hyperparameters above for adversarial training.

MNIST

Figure 5.3 (Top) shows the robustness of naturally trained models to L_∞- and L_2-adversarial attacks on MNIST. Against L_∞-adversarial attacks, RMSprop produce the most robust model with NAUC 0.33, followed by Adam (0.27), SGD with momentum (0.26), and SGD and Adagrad (0.22). Against L_2-adversarial attacks, SGD produces the most robust model with NAUC 0.49, followed by SGD with momentum (0.48), Adam (0.43), and Adagrad and RMSprop (0.41).

Against L_∞-adversarial attacks, RMSprop produces a model that appears qualitatively more robust than the next best performing model. This difference can be large at specific ϵ; for example at $\epsilon = 0.2$, the RMSprop model maintains a robust accuracy of 22%, while the Adam model has robust accuracy 7%. This is the only instance across all of our experiment in which we observe a notable qualitative difference between different algorithms.

Figure 5.3 (Bottom) shows the robustness of adversarially trained models. Training adversarially improves the robustness over naturally trained models regardless of the optimization algorithm and all optimization algorithms give qualitatively similar results. For L_∞-adversarial training, SGD produces the most robust model with NAUC 0.66, followed by SGD with momentum (0.65), Adam (0.64), RMSprop (0.63), and Adagrad (0.61). For L_2-adversarial training SGD with momentum and RMSprop produce models with NAUC 0.56, followed by SGD (0.54), Adam (0.53), and Adagrad (0.51). For adversarial training on MNIST, either SGD or SGD with momentum were among the top performers, with Adagrad always producing the worst performing model.

Adversarial training does seem to reduce the dependence on the choice of optimization algorithm, though does not completely remove it. Against L_∞-adversarial attacks at $\epsilon = 0.3$, the SGD model maintains robust accuracy of 91.5%, while the Adagrad model maintains robust accuracy 83.5%. We consider this difference noteworthy for MNIST. The difference is less pronounced for L_2-adversarial training.

SGD and SGD with momentum consistently outperform other optimization algorithms, yielding the best models in three out of four experiments. (Or one of the most robust models in the case of ties.) While the difference is qualitatively small, we believe that the consistency with which SGD or SGD with momentum produces the most robust model is noteworthy. Furthermore Adagrad seems to consistently under-perform other optimization algorithms. (Though this may depend on the domain [89].)

CIFAR10

Figure 5.4 (Top) shows the robustness of naturally trained models to L_∞- and L_2-adversarial attacks on CIFAR10. Against L_∞-adversarial attacks, SGD with momentum produces the most robust model with NAUC 0.06, followed by RMSprop (0.05), Adam and Adagrad (0.04), and SGD (0.02). Against L_2-adversarial attacks Adagrad produces the most robust models with NAUC 0.19, followed by SGD, Adam, and RMSprop (0.17), and SGD with momentum (0.15).

Figure 5.4 (Bottom) shows the robustness of adversarially trained models. For L_∞-adversarial training, SGD with momentum produced the most robust model with NAUC 0.26, followed by SGD (0.24), Adam (0.23), and Adagrad and RMSprop (0.22). For L_2-adversarial training, SGD

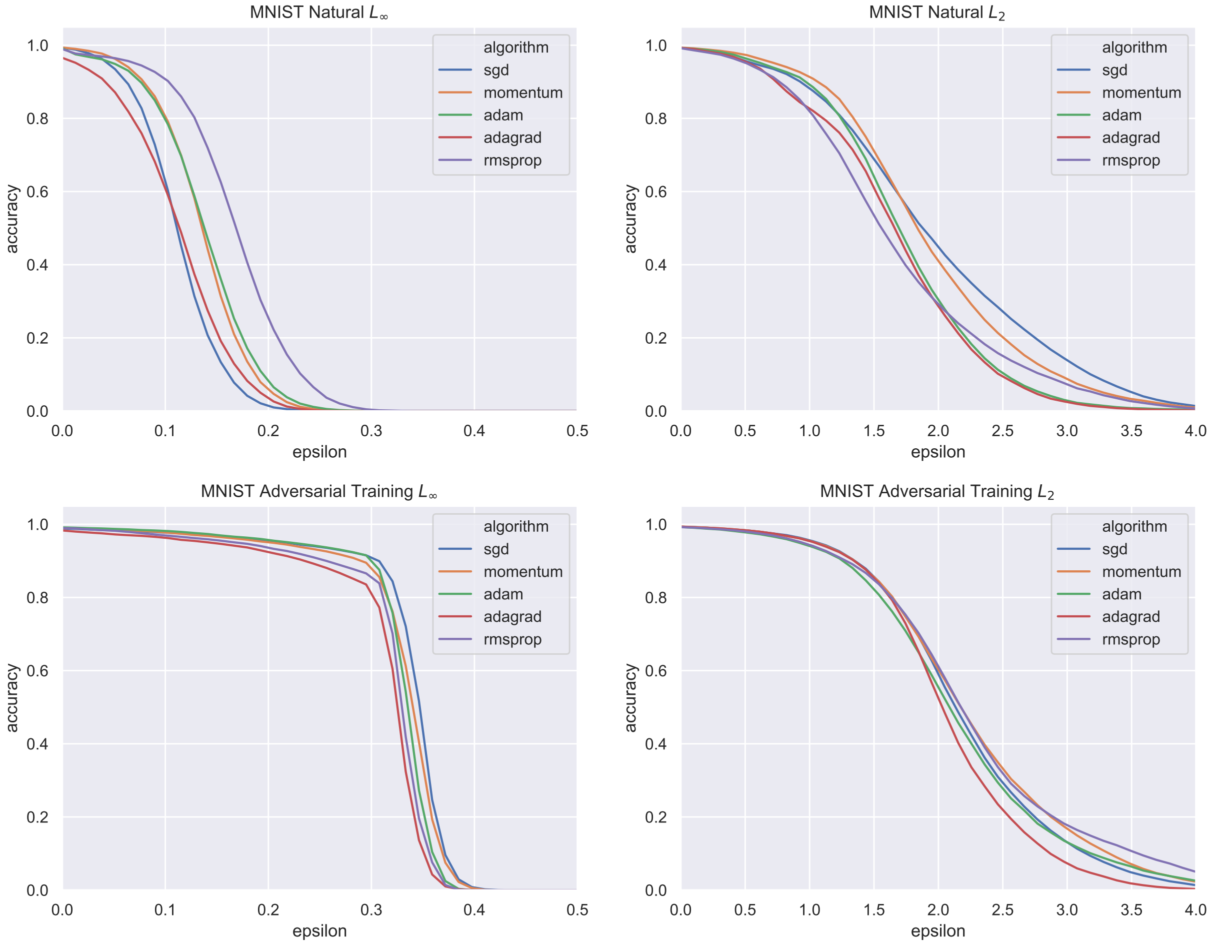

Figure 5.3: **Top**: Robust accuracy for naturally trained MNIST models against L_∞- and L_2-adversarial attacks. **Bottom**: Robust accuracy for adversarially trained MNIST models. Left, L_∞-adversarially trained models evaluated against L_∞-adversarial attacks; right L_2-adversarially trained models evaluated against L_2-adversarial attacks.

produces the most robust model with NAUC 0.51, followed by SGD with momentum (0.50), Adam and RMSprop (0.49), and Adagrad (0.47).

Adversarial training lessens the dependence on the choice of optimization algorithm. Against L_∞-adversarial attacks at $\epsilon = 0.031$ $(= 8/255)$, the SGD with momentum model maintains robust accuracy of 45%, while the RMSprop model maintains robust accuracy 34%. The difference is less pronounced for L_2-adversarial training.

SGD or SGD with momentum consistently outperform other optimization algorithms, yielding the best models in three out of four experiments. Again while the difference is qualitatively small, we believe that the consistency with which SGD or SGD with momentum produces the most robust model is noteworthy.

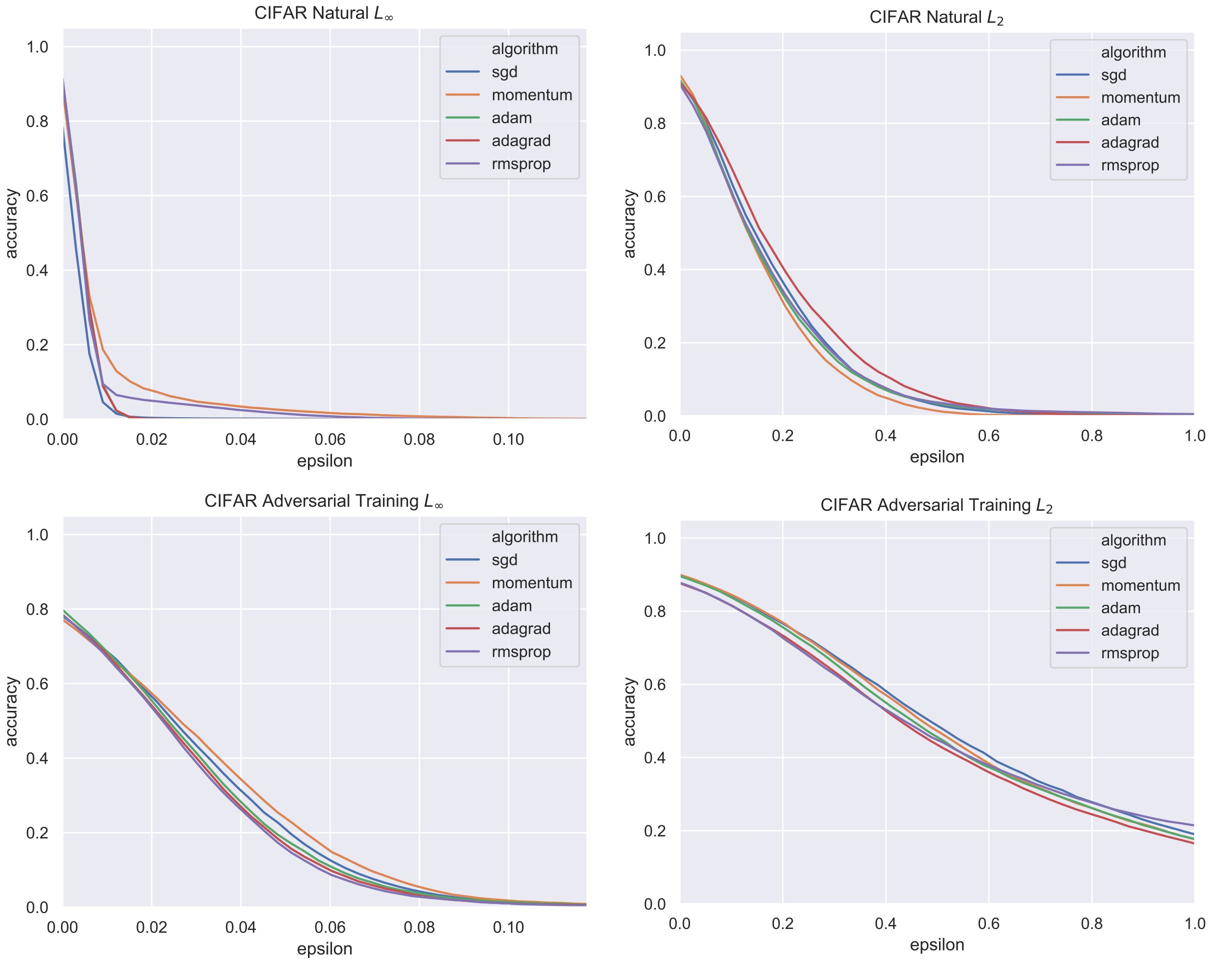

Figure 5.4: **Top**: Robust accuracy for naturally trained CIFAR10 models against L_∞- and L_2-adversarial attacks. **Bottom**: Robust accuracy for adversarially trained CIFAR10 models. Left, L_∞-adversarially trained models evaluated against L_∞-adversarial attacks; right L_2-adversarially trained models evaluated against L_2-adversarial attacks.

Bibliography

[1] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. "A Convergence Theory for Deep Learning via Over-Parameterization". In: *ICML*. 2019.

[2] Nina Amenta and Marshall W. Bern. "Surface Reconstruction by Voronoi Filtering". In: *Discrete & Computational Geometry* (1999).

[3] Nina Amenta, Marshall W. Bern, and David Eppstein. "The Crust and the beta-Skeleton: Combinatorial Curve Reconstruction". In: *Graphical Models and Image Processing* (1998).

[4] Nina Amenta, Sunghee Choi, and Ravi Kolluri. "The Power Crust". In: *Proceedings of the Sixth Symposium on Solid Modeling*. Association for Computing Machinery, 2001, pp. 249–260.

[5] Nina Amenta and Tamal Krishna Dey. *Normal Variation with Adaptive Feature Size*. 2007.

[6] Nina Amenta et al. "A Simple Algorithm for Homeomorphic Surface Reconstruction". In: *International Journal of Computational Geometry and Applications* (2002).

[7] Anish Athalye, Nicholas Carlini, and David A. Wagner. "Obfuscated Gradients Give a False Sense of Security: Circumventing Defenses to Adversarial Examples". In: *ICML*. 2018.

[8] Ivo Babuška and Abdul Kadir Aziz. "On the Angle Condition in the Finite Element Method". In: *SIAM Journal on Numerical Analysis* 13.2 (Apr. 1976), pp. 214–226.

[9] Harry Blum. "A transformation for extracting new descriptors of shape". In: *Models for Perception of Speech and Visual Forms* (1967).

[10] Jean-Daniel Boissonnat and Arijit Ghosh. "Manifold Reconstruction Using Tangential Delaunay Complexes". In: *Discrete & Computational Geometry* 51 (2014).

[11] Jean-Daniel Boissonnat and Steve Oudot. "Provably Good Sampling and Meshing of Surfaces". In: *Graphical Models* 67.5 (Sept. 2005), pp. 405–451.

[12] Jean-Daniel Boissonnat and Steve Oudot. "Provably Good Surface Sampling and Approximation". In: *Symposium on Geometry Processing*. Aachen, Germany: Eurographics Association, June 2003, pp. 9–18. ISBN: 1-58113-687-0.

[13] Jean-Daniel Boissonnat et al. "Meshing of Surfaces". In: *Effective Computational Geometry for Curves and Surfaces*. Ed. by Jean-Daniel Boissonnat and Monique Teillaud. Springer, 2006. Chap. 5, pp. 181–229.

[14] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[15] James H. Bramble and Miloš Zlámal. “Triangular Elements in the Finite Element Method”. In: *Mathematics of Computation* 24.112 (Oct. 1970), pp. 809–820.

[16] Sébastien Bubeck et al. “Adversarial examples from computational constraints”. In: *ICML*. 2019.

[17] Jianbo Chen, Michael I. Jordan, and Martin J. Wainwright. “HopSkipJumpAttack: Query-Efficient Decision-Based Adversarial Attack”. In: *CoRR* abs/1904.02144 (2019).

[18] Ho-Lun Cheng et al. “Dynamic Skin Triangulation”. In: *Discrete & Computational Geometry* 25.4 (Dec. 2001), pp. 525–568.

[19] Siu-Wing Cheng, Tamal K. Dey, and Edgar A. Ramos. “Manifold reconstruction from point samples”. In: *Proceedings of the Symposium on Discrete Algorithms (SODA)*. 2005.

[20] Siu-Wing Cheng, Tamal K. Dey, and Jonathan Richard Shewchuk. *Delaunay Mesh Generation*. Boca Raton, Florida: CRC Press, Dec. 2012.

[21] Siu-Wing Cheng, Tamal Krishna Dey, and Edgar A. Ramos. “Delaunay Refinement for Piecewise Smooth Complexes”. In: *Discrete & Computational Geometry* 43.1 (2010), pp. 121–166.

[22] Siu-Wing Cheng et al. “Sliver exudation”. In: *Journal of the ACM* 47 (2000).

[23] L. Paul Chew. “Constrained Delaunay Triangulations”. In: *Algorithmica* 4.1 (1989), pp. 97–108.

[24] Felipe Codevilla et al. “End-to-end Driving via Conditional Imitation Learning”. In: *ICRA*. 2018.

[25] Jeremy M. Cohen, Elan Rosenfeld, and J. Zico Kolter. “Certified Adversarial Robustness via Randomized Smoothing”. In: *ICML*. 2019.

[26] Corinna Cortes and Vladimir Vapnik. “Support-Vector Networks”. In: *Machine Learning* 20 (1995).

[27] Akshay Degwekar, Preetum Nakkiran, and Vinod Vaikuntanathan. “Computational Limitations in Robust Classification and Win-Win Results”. In: *COLT*. 2019.

[28] Tamal K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge University Press, 2007.

[29] Tamal K. Dey and Samrat Goswami. “Provable surface reconstruction from noisy samples”. In: *Proceedings of the Symposium on Computational Geometry (SoCG)*. 2004.

[30] Tamal K. Dey and Piyush Kumar. “A Simple Provable Algorithm for Curve Reconstruction”. In: *Proceedings of the Symposium on Discrete Algorithms (SODA)*. 1999.

[31] Tamal K. Dey et al. “Critical Points of Distance to an ϵ-sampling of a Surface and Flow-Complex-Based Surface Reconstruction”. In: *International Journal of Computational Geometry and Applications* (2008).

[32] Tamal Krishna Dey and Joshua A. Levine. “Delaunay Meshing of Isosurfaces”. In: *Visual Computer* 24.6 (June 2008), pp. 411–422.

[33] Gavin Weiguang Ding et al. “Max-Margin Adversarial (MMA) Training: Direct Input Space Margin Maximization through Adversarial Training”. In: *CoRR* abs/1812.02637 (2018).

[34] Simon S. Du et al. “Gradient Descent Finds Global Minima of Deep Neural Networks”. In: *ICML*. 2019.

[35] John C. Duchi, Elad Hazan, and Yoram Singer. “Adaptive Subgradient Methods for Online Learning and Stochastic Optimization”. In: *Journal of Machine Learning Research* (2011).

[36] Herbert Edelsbrunner and Nimish R. Shah. “Triangulating Topological Spaces”. In: *International Journal of Computational Geometry and Applications* (Aug. 1997).

[37] Gamaleldin F. Elsayed et al. “Large Margin Deep Networks for Classification”. In: *CoRR* abs/1803.05598 (2018). arXiv: `1803.05598`. URL: `http://arxiv.org/abs/1803.05598`.

[38] Andre Esteva et al. “Dermatologist-level classification of skin cancer with deep neural networks”. In: *Nature* (2017).

[39] Timon Gehr et al. “AI2: Safety and robustness certification of neural networks with abstract interpretation”. In: *IEEE Symposium on Security and Privacy*. 2018.

[40] Justin Gilmer et al. “Adversarial Spheres”. In: *CoRR* abs/1801.02774 (2018). arXiv: `1801.02774`. URL: `http://arxiv.org/abs/1801.02774`.

[41] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. “Explaining and Harnessing Adversarial Examples”. In: *ICLR*. 2014.

[42] Suriya Gunasekar et al. “Characterizing Implicit Bias in Terms of Optimization Geometry”. In: *ICML*. 2018.

[43] Suriya Gunasekar et al. “Implicit Bias of Gradient Descent on Linear Convolutional Networks”. In: *NeurIPS*. 2018.

[44] R. H. Hardin, N. J. A. Sloane, and W. D. Smith. *Coverings by Points on a Sphere*. Web page. Feb. 1994.

[45] Kaiming He et al. “Deep Residual Learning for Image Recognition”. In: *CVPR*. 2016.

[46] Dan Hendrycks et al. “Natural Adversarial Examples”. In: *CoRR* abs/1907.07174 (2019).

[47] Judy Hoffman, Daniel A. Roberts, and Sho Yaida. “Robust Learning with Jacobian Regularization”. In: *CoRR* abs/1908.02729 (2019).

[48] Andrew Ilyas et al. “Adversarial Examples Are Not Bugs, They Are Features”. In: *CoRR* abs/1905.02175 (2019).

[49] Russell Impagliazzo. “A Personal View of Average-Case Complexity”. In: *Proceedings of the Tenth Annual Structure in Complexity Theory Conference*. 1995.

[50] Daniel Jakubovitz and Raja Giryes. "Improving DNN Robustness to Adversarial Attacks Using Jacobian Regularization". In: *ECCV*. 2018.

[51] Ajil Jalal et al. "The Robust Manifold Defense: Adversarial Training using Generative Models". In: *CoRR* abs/1712.09196 (2019).

[52] Camille Jordan. "Essai sur la Géométrie à *n* Dimensions". In: *Bulletin de la Société Mathématique de France* 3 (1875), pp. 103–174.

[53] Daniel Kang et al. "Testing Robustness Against Unforeseen Adversaries". In: *arXiv preprint arXiv:1908.08016* (2019).

[54] Andrej Karparthy. *A peak at trends in machine learning*. `https://medium.com/@karpathy/a-peek-at-trends-in-machine-learning-ab8a1085a106`. 2017.

[55] Marc Khoury and Dylan Hadfield-Menell. "Adversarial Training with Voronoi Constraints". In: *CoRR* abs/1905.01019 (2019).

[56] Marc Khoury and Dylan Hadfield-Menell. "On the Geometry of Adversarial Examples". In: *CoRR* abs/1811.00525 (2018).

[57] Marc Khoury and Jonathan Richard Shewchuk. "Approximation Bounds for Interpolation and Normals on Triangulated Surfaces and Manifolds". In: *CoRR* abs/1911.03424 (2019). URL: `http://arxiv.org/abs/1911.03424`.

[58] Marc Khoury and Jonathan Richard Shewchuk. "Fixed Points of the Restricted Delaunay Triangulation Operator". In: *Proceedings of the Symposium on Computational Geometry (SoCG)*. 2016.

[59] Diederik Kingma and Jimmy Ba. "Adam: A Method for Stochastic Optimization". In: *ICLR*. 2015.

[60] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. "ImageNet Classification with Deep Convolutional Neural Networks". In: *NIPS*. 2012.

[61] Mathias Lecuyer et al. "Certified robustness to adversarial examples with differential privacy". In: *CoRR* abs/1802.03471 (2018).

[62] Der-Tsai Lee and Arthur K. Lin. "Generalized Delaunay Triangulations for Planar Graphs". In: *Discrete & Computational Geometry* 1 (1986), pp. 201–217.

[63] Sergey Levine, Nolan Wagener, and Pieter Abbeel. "Learning contact-rich manipulation skills with guided policy search". In: *ICRA*. 2015.

[64] Xuezhi Liang et al. "Soft-Margin Softmax for Deep Classification". In: *ICONIP*. 2017.

[65] Weiyang Liu et al. "Large-Margin Softmax Loss for Convolutional Neural Networks". In: *ICML*. 2016.

[66] Aleksander Madry et al. "Towards Deep Learning Models Resistant to Adversarial Attacks". In: *ICLR*. 2018.

[67] Matthew Mirman, Timon Gehr, and Martin T. Vechev. “Differentiable Abstract Interpretation for Provably Robust Neural Networks”. In: *ICML*. 2018.

[68] Kamil Nar et al. “Cross-Entropy Loss and Low-Rank Features Have Responsibility for Adversarial Examples”. In: *CoRR* abs/1901.08360 (2019).

[69] Steve Oudot, Laurent Rineau, and Mariette Yvinec. “Meshing Volumes Bounded by Smooth Surfaces”. In: *Proceedings of the 14th International Meshing Roundtable*. San Diego, California: Springer, Sept. 2005, pp. 203–219.

[70] Andrew Pressley. *Elementary Differential Geometry*. Springer Science & Business Media, 2010.

[71] Maithra Raghu et al. “On the Expressive Power of Deep Neural Networks”. In: *ICML*. 2017.

[72] Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. “Certified Defenses against Adversarial Examples”. In: *ICLR*. 2018.

[73] V. T. Rajan. “Optimality of the Delaunay Triangulation in $\mathbb{R}^d$”. In: *Proceedings of the Seventh Annual Symposium on Computational Geometry*. North Conway, New Hampshire, June 1991, pp. 357–363.

[74] Laurent Rineau and Mariette Yvinec. “Meshing 3D Domains Bounded by Piecewise Smooth Surfaces”. In: *Proceedings of the 16th International Meshing Roundtable*. Seattle, Washington: Springer, Oct. 2007, pp. 443–460.

[75] Hadi Salman et al. “Provably Robust Deep Learning via Adversarially Trained Smoothed Classifiers”. In: *NeurIPS*. 2019.

[76] Ludwig Schmidt et al. “Adversarially Robust Generalization Requires More Data”. In: *NIPS*. 2018.

[77] E. Schönhardt. “Über die Zerlegung von Dreieckspolyedern in Tetraeder”. In: *Mathematische Annalen* 98 (1928), pp. 309–312.

[78] Lukas Schott et al. “Towards the first adversarially robust neural network model on MNIST”. In: *CoRR* abs/1805.09190 (2018). arXiv: `1805.09190`. URL: `https://arxiv.org/abs/1805.09190`.

[79] Raimund Seidel. “Constrained Delaunay Triangulations and Voronoi Diagrams with Obstacles”. In: *1978–1988 Ten Years IIG*. Ed. by H. S. Poingratz and W. Schinnerl. Institute for Information Processing, Graz University of Technology, 1988, pp. 178–191.

[80] Ali Shafahi et al. “Are adversarial examples inevitable?” In: *ICLR*. 2019.

[81] John Shawe-Taylor and Nello Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.

[82] Jonathan Richard Shewchuk. “General-Dimensional Constrained Delaunay Triangulations and Constrained Regular Triangulations, I: Combinatorial Properties”. In: *Discrete & Computational Geometry* 39.1–3 (Mar. 2008), pp. 580–637.

[83] Jonathan Richard Shewchuk. “What Is a Good Linear Element? Interpolation, Conditioning, and Quality Measures”. In: *Proceedings of the 11th International Meshing Roundtable*. Ithaca, New York: Sandia National Laboratories, Sept. 2002, pp. 115–126.

[84] Gagandeep Singh et al. “An abstract domain for certifying neural networks”. In: *PACMPL* (2019).

[85] Aman Sinha, Hongseok Namkoong, and John Duchi. “Certifying Some Distributional Robustness with Principled Adversarial Training”. In: *ICLR*. 2018.

[86] Shizhao Sun et al. “On the Depth of Deep Neural Networks: A Theoretical View”. In: *AAAI*. 2016.

[87] Christian Szegedy et al. “Intriguing properties of neural networks”. In: *CoRR* abs/1312.6199 (2013). URL: `http://arxiv.org/abs/1312.6199`.

[88] Tijmen Tieleman and Geoffrey Hinton. *RMSprop: Divide the gradient by a running average of its recent magnitude*. 2012.

[89] Alexandru Tifrea, Gary Bécigneul, and Octavian-Eugen Ganea. “Poincare Glove: Hyperbolic Word Embeddings”. In: *ICLR*. 2019.

[90] Dimitris Tsipras et al. “Robustness May Be at Odds with Accuracy”. In: *ICLR*. 2019.

[91] Yizhen Wang, Somesh Jha, and Kamalika Chaudhuri. “Analyzing the Robustness of Nearest Neighbors to Adversarial Examples”. In: *ICML*. 2018.

[92] Tsui-Wei Weng et al. “Towards Fast Computation of Certified Robustness for ReLU Networks”. In: *ICML*. 2018.

[93] Ashia C. Wilson et al. “The Marginal Value of Adaptive Gradient Methods in Machine Learning”. In: *NIPS*. 2017.

[94] Eric Wong and J. Zico Kolter. “Provable Defenses against Adversarial Examples via the Convex Outer Adversarial Polytope”. In: *ICML*. 2018.

[95] Eric Wong et al. “Scaling provable adversarial defenses”. In: *NeurIPS*. 2018.

[96] Yonghui Wu et al. “Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation”. In: *CoRR* abs/1609.08144 (2016). arXiv: `1609.08144`. URL: `http://arxiv.org/abs/1609.08144`.

[97] Hongyang Zhang et al. “Theoretically Principled Trade-off between Robustness and Accuracy”. In: *ICML*. 2019.

Appendix A

Proofs for Chapter 3

A.1 Useful Known Results

Lemma 38 (Feature Translation Lemma)**.** *Let* $\Sigma \subset \mathbb{R}^3$ *be a smooth surface and let* $p, q \in \Sigma$ *be points on* Σ *such that* $|pq| \leq \epsilon \operatorname{lfs}(p)$ *for some* $\epsilon < 1$. *Then*

$$\operatorname{lfs}(p) \leq \frac{1}{1-\epsilon}\operatorname{lfs}(q) \quad \textit{and} \quad |pq| \leq \frac{\epsilon}{1-\epsilon}\operatorname{lfs}(q).$$

Proof. By the definition of the local feature size, there is a medial axis point m such that $|qm| = \operatorname{lfs}(q)$. By the Triangle Inequality, $\operatorname{lfs}(p) \leq |pm| \leq |pq| + |qm| \leq \epsilon \operatorname{lfs}(p) + \operatorname{lfs}(q)$. Rearranging terms gives $\operatorname{lfs}(p) \leq \operatorname{lfs}(q)/(1-\epsilon)$. The second claim follows immediately. □

Lemma 39 (Feature Ball Lemma [28])**.** *If a geometric closed d-ball B intersects a k-manifold* $\Sigma \subset \mathbb{R}^d$ *without boundary at more than one point and either (i)* $B \cap \Sigma$ *is not a topological k-ball or (ii)* $\partial(B \cap \Sigma)$ *is not a topological* $(k-1)$*-sphere, then B contains a medial axis point.*

Lemma 40. *Let s be a line segment with endpoints* $p, q \in \Sigma$ *where* $\Sigma \subset \mathbb{R}^d$ *is a k-manifold without boundary. Let* B_s *be the* diametric ball *of s, the smallest closed d-ball such that* $B_s \supset s$ *(for which s is a diameter of* B_s*). If* $d(p, q) \leq \rho \operatorname{lfs}(p)$ *for some* $\rho < 1$*, then* $B_s \cap \Sigma$ *is a topological* $(k-1)$*-ball.*

Proof. Let B_s be the diametric ball of s. Suppose that $B_s \cap \Sigma$ is not a topological 2-ball. As $B_s \cap \Sigma$ contains more than one point (p and q), by the Feature Ball Lemma (Lemma 39), there exists a medial axis point $m \in B_s$. This implies that $\operatorname{lfs}(p) \leq d(p, m) \leq d(p, q) \leq \rho \operatorname{lfs}(p)$, which is a contradiction. □

Lemma 41 (Small Circumradius Lemma [28])**.** *Let V be an* ϵ*-sample of a smooth surface* $\Sigma \subset \mathbb{R}^3$ *for* $\epsilon < 1$. *Let* $\tau \in \operatorname{Del}|_\Sigma V$ *be a restricted Delaunay triangle and let* $x = \tau^* \cap \Sigma$ *be the intersection of* τ*'s dual Voronoi edge with* Σ. *Then the circumradius of* τ *is at most* $\epsilon \operatorname{lfs}(x)$.

A.2 Extended Voronoi Cell Boundaries

There are only three phenomena that can determine the boundary of an extended Voronoi cell. (1) Portions of a cell's boundary may be determined by hyperplanes, each hyperplane being equidistant from two sites. For example, a point on the boundaries of two cells $\mathrm{Vor}|_{\tilde{\Sigma}_S} v$ and $\mathrm{Vor}|_{\tilde{\Sigma}_S} w$ might lie on the hyperplane that orthogonally bisects the line segment vw. (2) A cell's boundary may include some of the rays that bound the extrusions. The boundary of $\mathrm{Vor}|_{\tilde{\Sigma}_S} v$ includes two such rays for each segment in S that adjoins v. (3) Portions of a cell's boundary may be determined by a shadow cast by a portal. For example, consider a point $x \in \mathrm{Vor}|_{\tilde{\Sigma}_S} v$ such that the line segment vx intersects the boundary of a portal. The boundary of a portal does not block visibility, but the relative interior of a portal does block visibility in the principal branch. Suppose that the portal hides some points arbitrarily close to x (i.e., some points in every open neighborhood of x on Σ) so they are not visible from v. Typically, each of those shadowed points belongs to one or more other extended Voronoi cells (even if v is the closest site), so x can be on the boundary of a cell $\mathrm{Vor}|_{\tilde{\Sigma}_S} w$ without being in $\mathrm{Vor}|_{\tilde{\Sigma}_S} w$. In that case, w's cell is not a closed point set and portal shadows are responsible for the open portions of its boundary.

In this section, we establish that under suitable sampling conditions, the third phenomenon cannot happen, so the boundaries of all the extended Voronoi cells are determined solely by bisecting hyperplanes and extrusion boundary rays, and all the extended Voronoi cells are closed point sets. The following theorem states those sampling conditions.

Theorem 42 (Shadow Theorem)**.** *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface without boundary. Define a set of segments S and a set of portal curves Z as described in Section 3.5 such that for every segment $s = pq \in S$, $d(p,q) \leq 0.47\,\mathrm{lfs}(p)$. Let $V \subset \Sigma$ be a constrained ϵ-sample of (Σ, S, Z) for some $\epsilon \leq 1$.*

Then for every site $v \in V$ and every point x in the extended Voronoi cell $\mathrm{Vor}|_{\tilde{\Sigma}} v$, the relative interior of the line segment xv does not intersect the boundary of a portal. Hence, for every site $v \in V$ and every point x on the boundary of $\mathrm{Vor}|_{\tilde{\Sigma}} v$, x lies on a hyperplane that bisects a line segment vw for some $w \in V$ or on a ray that bounds an extrusion.

Proof. Follows from Lemmas 44 and 46 below. □

The proof of the Shadow Theorem divides into two parts: we prove it first for the case where x is in the principal branch (the Principal Shadow Lemma, Lemma 44), then for the case where x lies on an extrusion (the Extrusion Shadow Lemma, Lemma 46).

Lemma 43. *Let $\Sigma \subset \mathbb{R}^3$ be a point set and let M be its medial axis. Consider three points $x \in \mathbb{R}^3$ and $y, z \in \mathbb{R}^3 \setminus M$. Suppose that $\tilde{y} \neq \tilde{z}$ and y lies on the line segment xz. Then $d(x, \tilde{y}) < d(x, \tilde{z})$.*

Proof. Let Π be the plane that bisects the line segment $\tilde{y}\tilde{z}$. As $\tilde{y}$ is the unique point nearest y on Σ, $d(y, \tilde{y}) < d(y, \tilde{z})$, so y lies on the same side of Π as $\tilde{y}$. Symmetrically, z lies on the same side of Π as $\tilde{z}$. As $y \in xz$, x lies on the same side of Π as $\tilde{y}$. Therefore, $d(x, \tilde{y}) < d(x, \tilde{z})$. □

Lemma 44 (Principal Shadow Lemma)**.** *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface without boundary. Let $V \subset \Sigma$ be a finite set of sites. Consider a site $v \in V$ and a point x in v's principal Voronoi cell*

$\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, v$. *If* $d(x, v) \leq \mathrm{lfs}(x)$, *then the relative interior of the line segment* xv *does not intersect the boundary of a portal.*

Proof. Suppose for the sake of contradiction that the relative interior of xv intersects one or more portal boundaries. Let y_1 be the intersection point closest to x. Observe that x is visible from y_1 (because x is visible from v). As $d(x, y_1) < d(x, v) \leq \mathrm{lfs}(x)$, y_1 does not lie on the medial axis. But y_1 lies on a portal boundary, so y_1 must lie on the normal segment of a site in V, and that site is located at the point $\tilde{y}_1 \neq v$. By Lemma 43, $d(x, \tilde{y}_1) < d(x, v)$.

Although x is closer to the site at $\tilde{y}_1$ than it is to v, x is in v's principal Voronoi cell rather than $\tilde{y}_1$'s cell, so x is not visible from the site at $\tilde{y}_1$, even though x is visible from y_1. Let z_1 be the point nearest $\tilde{y}_1$ on the line segment $y_1\tilde{y}_1$ that can see x. Observe that $\tilde{z}_1 = \tilde{y}_1$ (all three points y_1, $\tilde{y}_1$, and z_1 lie on the normal segment of $\tilde{y}_1$). The relative interior of the line segment xz_1 intersects one or more portal boundaries (because as you slide from y_1 to $\tilde{y}_1$, z_1 is the last point that can see x). Let y_2 be the intersection point closest to x. Then x is visible from y_2, and y_2 lies on the boundary of a portal.

As $d(x, y_2) < d(x, z_1) \leq \max\{d(x, y_1), d(x, \tilde{y}_1)\} < d(x, v) \leq \mathrm{lfs}(x)$, y_2 does not lie on the medial axis. But y_2 lies on a portal boundary, so y_2 must lie on the normal segment of a site in V, and that site is located at the point $\tilde{y}_2 \neq \tilde{z}_1$. By Lemma 43, $d(x, \tilde{y}_2) < d(x, \tilde{z}_1)$; hence $d(x, \tilde{y}_2) < d(x, \tilde{y}_1) < d(x, v)$.

As $x \in \mathrm{Vor}\,|_{\overline{\Sigma}_S}\, v$ (rather than $\tilde{y}_2$'s cell), x is not visible from the site at $\tilde{y}_2$, even though x is visible from y_2. By inductively repeating the argument we can construct an infinite sequence of sites not visible from x such that each successive site is closer to x than the previous site. But V contains finitely many sites. The result follows by contradiction. □

Next, we prove another Shadow Lemma for any point x on an extrusion. The proof of the Extrusion Shadow Lemma is almost the same as the proof of the Principal Shadow Lemma, but there is one major complication: that proof uses the fact that no point in the sequence $y_1, y_2, \ldots$ lies on the medial axis. This fact is more difficult to prove when x lies on an extrusion (rather than on Σ).

Consider a segment $s \in S$ with endpoints $p, q \in V$ such that $d(p, q) \leq \rho\, \mathrm{lfs}(p)$ for some $\rho \leq 0.47$. Let F be the open ball with center p and radius $\mathrm{lfs}(p)$. By the definition of lfs, F does not intersect the medial axis of Σ. Let c be the midpoint of s. Let B_λ be the closed ball with center c and radius $\lambda\, \mathrm{lfs}(p)$, where λ is the function of ρ defined in Theorem 20. Observe that $B_\lambda \subset F$, as the distance from p to any point in B_λ is at most $(\rho/2 + \lambda)\, \mathrm{lfs}(p) < 0.705\, \mathrm{lfs}(p)$.

Let h_s be the cutting plane for s, let $\zeta_s \subset h_s \cap \Sigma$ be the portal curve for s, and let b_s be a unit vector normal to h_s. Let Σ_s^+ be the extrusion of s's portal curve in the direction b_s. Recall that by Theorem 20, every site whose extended Voronoi cell intersects Σ_s^+ must lie in B_λ.

Let W be the set of all points $\{x + \omega b_s : x \in B_\lambda \text{ and } \omega \geq 0\}$; that is, every point that is in B_λ or in the direction b_s from a point in B_λ. The set W has the shape of a wiener that is infinite in one direction. The premise of W is that it is a convex point set that encloses both the ball B_λ and the extrusion Σ_s^+. However, W is defined with respect to the space $\mathbb{R}^3$ whereas Σ_s^+ is embedded in a secondary branch of $\widetilde{X}$. If Σ_s^+ were in $\mathbb{R}^3$ (but the point coordinates were unchanged), we could write $\Sigma_s^+ \subset W$.

Lemma 45. *Define s, p, q, F, B_λ, h_s, and W as above. Let h_s^- be the closed halfspace on the side of h_s opposite from Σ_s^+. Consider a point $z \in \Sigma \cap B_\lambda \cap h_s^-$ and let ℓ_z be its normal segment. Then $W \cap \ell_z \subset F$.*

Proof. Suppose for the sake of contradiction that there is a point $u \in W \cap \ell_z$ such that $u \notin F$. Then $d(p, u) \geq \operatorname{lfs}(p)$ and $u \notin B_\lambda$. Let B_u be the open ball with center u whose boundary passes through z. As u lies on z's normal segment, B_u does not intersect Σ.

Let B and B' be the two open balls of radius $\operatorname{lfs}(p)$ tangent to Σ at p; neither ball intersects Σ. As Σ is a surface without boundary that passes through p, it partitions $\mathbb{R}^3$ into two disjoint components, with B included in one and B' in the other. One component must include B_u too; suppose without loss of generality that B' is in the same component as B_u. It follows that B_u is disjoint from B.

Let o and o' be the centers of B and B', respectively. As B_u and B are disjoint, $d(u, o) \geq d(u, z) + \operatorname{lfs}(p)$. Let $u_\perp$ denote the distance from u to the cutting plane h_s. Whereas $z \in h_s^-$, $u \in W \setminus F$ must lie on the positive side of h_s (opposite from h_s^-). Therefore, $d(u, z) \geq u_\perp$ and $d(u, o) \geq u_\perp + \operatorname{lfs}(p)$.

As the cutting plane h_s is parallel to p's normal, $o, o' \in h_s$. Let $\bar{u}$ be the orthogonal projection of u onto h_s. By Pythagoras' Theorem, $d(u, o)^2 = u_\perp^2 + d(\bar{u}, o)^2$. Combining this with the last inequality gives $d(\bar{u}, o)^2 \geq 2u_\perp \operatorname{lfs}(p) + \operatorname{lfs}(p)^2$. Therefore, $u_\perp \leq (d(\bar{u}, o)^2/\operatorname{lfs}(p) - \operatorname{lfs}(p))/2$.

The fact that $u \in W$ implies that $\bar{u} \in B_\lambda$, so we can write $d(\bar{u}, o) \leq d(\bar{u}, c) + d(c, o) \leq \lambda \operatorname{lfs}(p) + d(c, o)$. To find an upper bound for $d(c, o)$, consider a coordinate system that places the site p at the origin, the cutting plane h_s in the x-y plane, the point o at the coordinate $(0, \operatorname{lfs}(p), 0)$, and the point o' at the coordinate $(0, -\operatorname{lfs}(p), 0)$. Then we write $q = (q_x, q_y, 0)$ and $c = (q_x/2, q_y/2, 0)$. With this notation, $d(q, o')^2 = q_x^2 + (q_y + \operatorname{lfs}(p))^2 = \|q\|^2 + 2\operatorname{lfs}(p)q_y + \operatorname{lfs}(p)^2$ and $d(c, o)^2 = (q_x/2)^2 + (q_y/2 - \operatorname{lfs}(p))^2 = \|q\|^2/4 - \operatorname{lfs}(p)q_y + \operatorname{lfs}(p)^2$. Adding half the first equation to the second (to eliminate q_y) gives $d(q, o')^2/2 + d(c, o)^2 = 3\|q\|^2/4 + 3\operatorname{lfs}(p)^2/2$. Observe that $\|q\| = d(p, q) \leq \lambda \operatorname{lfs}(p)$. As $q \notin B'$ (because $q \in \Sigma$), $d(q, o') \geq \operatorname{lfs}(p)$. Hence $d(c, o)^2 \leq (3\lambda^2/4 + 1)\operatorname{lfs}(p)^2$.

Therefore, $d(\bar{u}, o) \leq (\lambda + \sqrt{3\lambda^2/4 + 1})\operatorname{lfs}(p) < 1.55 \operatorname{lfs}(p)$ and $u_\perp < 0.702 \operatorname{lfs}(p)$. Thus $d(p, u)^2 = d(p, \bar{u})^2 + u_\perp^2 \leq (d(p, c) + d(c, \bar{u}))^2 + u_\perp^2 < (\rho \operatorname{lfs}(p)/2 + \lambda \operatorname{lfs}(p))^2 + 0.702^2 \operatorname{lfs}(p)^2 < 0.99 \operatorname{lfs}(p)^2$. But this contradicts the fact that $u \notin F$. Hence there is no point $u \in W \cap \ell_z$ such that $u \notin F$. □

Lemma 46 (Extrusion Shadow Lemma)**.** *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface without boundary. Let $V \subset \Sigma$ be a finite set of sites. Let $s \in S$ be a segment with endpoints $p, q \in V$, and suppose that $d(p, q) \leq 0.47 \operatorname{lfs}(p)$. Consider a site $v \in V$ and a point x in v's extended Voronoi cell $\operatorname{Vor}|_{\widetilde{\Sigma}} v$ such that x lies on the extrusion Σ_s^+ (or Σ_s^-). Consider the line segment $xv \subset \widetilde{X}$. The relative interior of xv does not intersect the boundary of a portal (including the boundary of P_s).*

To clarify the interpretation of Lemma 46, note that xv lies partly in a secondary space and partly in the principal branch. The portion of xv that is solely in the secondary space cannot intersect any portal (regardless of matching point coordinates). The lemma focuses on the portion in the principal branch.

Proof. The proof is identical to that of Lemma 44, except that we employ an entirely different argument to establish that no point in the sequence $y_1, y_2, \ldots$ lies on the medial axis.

If v is a vertex of s then the result follows immediately, so assume that $v \in V \setminus \{p, q\}$. Define F, B_λ, h_s, ζ_s, and W as in the preamble before Lemma 45. Let h_s^- be the closed halfspace on the side of h_s opposite from Σ_s^+. By Theorem 18, $v \in h_s^-$. By Theorem 20, $v \in B_\lambda \subset W \cap F$.

Although the wiener W is defined in the Euclidean space $\mathbb{R}^3$, we say that W *encloses* a point set A if for every point $a \in A$ there is a point in W with the same coordinates as a, regardless of whether a is in the principal branch or a secondary branch. Observe that W encloses Σ_s^+ and B_λ. As $x \in \Sigma_s^+$, $v \in B_\lambda$, and W is convex, it follows that W encloses xv.

Suppose for the sake of contradiction that the relative interior of xv intersects one or more portal boundaries. Let y_1 be the intersection point closest to x. Let $u = xv \cap P_s$ be the point in the principal branch where the line segment xv exits the portal P_s. (If there is more than one such point, let u be the one nearest x. Usually uv lies entirely in the principal branch, but sometimes it does not because it intersects P_s or another portal in multiple points.) As y_1 is in the principal branch, $y_1 \in uv$. Observe that W encloses y_1 and u, and u lies on the normal segment $\ell_{\tilde{u}}$ of a point $\tilde{u}$ on the portal curve ζ_s. By Lemma 45, $W \cap \ell_{\tilde{u}} \subset F$, so $u \in F$. Recall that F is a medial-free open ball and $v \in F$, so F encloses uv and $y_1 \in F$. This confirms that y_1 is not a medial axis point.

As y_1 is on a portal boundary but not on the medial axis, there is a site at $\tilde{y}_1$. We claim that $\tilde{y}_1 \in h_s^-$; suppose for the sake of contradiction that $\tilde{y}_1 \notin h_s^-$. Observe that as uv does not intersect the medial axis, the nearest point map ν is continuous over uv and $\nu(uv)$ is a connected path on Σ. As $\tilde{y}_1$ is in the principal branch and lies on $\nu(uv)$, the path $\nu(uy_1)$ must somewhere (on the way from $\tilde{u} \in \zeta_s$ to $\tilde{y}_1$) exit h_s^-, without entering the portal P_s, to reach $\tilde{y}_1$. Let $w \in \nu(uy_1) \cap h_s$ be a point where the path crosses h_s without entering P_s. Hence $w \in \Sigma \cap h_s$ but w is not in the relative interior of the portal curve ζ_s. Let $\bar{x}$ be the orthogonal projection of x onto h_s, and recall that $\bar{x} \in \zeta_s$. As w is not in the relative interior of ζ_s, $d(\bar{x}, w) \geq \min\{d(\bar{x}, p), d(\bar{x}, q)\}$. As $x\bar{x}$ is orthogonal to h_s and $\bar{x}, w, p, q \in h_s$, we also have $d(x, w) \geq \min\{d(x, p), d(x, q)\}$. But by Lemma 43, $d(x, w) < d(x, v)$; and as $x \in \mathrm{Vor}|_\Sigma v$, $d(x, w) < d(x, v) \leq \min\{d(x, p), d(x, q)\}$. This is a contradiction; hence $\tilde{y}_1 \in h_s^-$ as claimed.

As $\tilde{y}_1 \in h_s^-$, by Theorem 20, $\tilde{y}_1 \in B_\lambda \subset W \cap F$. As $y_1 \in W \cap F$, $\tilde{y}_1 \in W \cap F$, and W and F are convex, it follows that the point $z_1 \in y_1\tilde{y}_1$ discussed in the proof of Lemma 44 is also in $W \cap F$. By repeating the argument of the previous two paragraphs with v replaced by z_1, we show that $y_2 \in W \cap F$, and hence y_2 is not a medial axis point. We repeat the argument inductively for y_3, y_4, etc. The rest of the proof proceeds as in the proof of Lemma 44. □

A.3 The Nearest Point Map Is a Homeomorphism

Here we prove that for a sufficiently dense sample V of a smooth surface $\Sigma \subset \mathbb{R}^3$ without boundary, with a suitable encroachment condition, the nearest point map (restricted to the restricted CDT) is a homeomorphism from the underlying space of the restricted CDT $\mathrm{Del}|_\Sigma V$ to Σ. (Recall that the nearest point map ν maps any point $x \in \mathbb{R}^3 \setminus M$ to the point $\nu(x)$ nearest x on Σ. In this section, we use the abbreviation $\tilde{x}$ to denote $\nu(x)$.) Specifically, we suppose V is a constrained 0.3202-sample of (Σ, S, H) and the encroachment condition described in Section 3.5 holds. We note that in the unconstrained case (i.e., for classical restricted Delaunay triangulations), our sampling constant is substantially better than those in the classical proofs from the literature on provably good surface reconstruction [28]: we prove homeomorphism for a 0.3202-sample instead of merely for a 0.08-sample. This reduces the number of samples required by a factor of about 16 (the square of 0.32/0.08).

Unlike the classical proofs that the restricted Delaunay triangulation is homeomorphic to the underlying manifold, our proof does not use the Topological Ball Theorem of Edelsbrunner and Shah [36]. The Topological Ball Theorem cannot be applied to the restricted *constrained* Delaunay triangulation because it depends on the barycentric subdivision of the Delaunay triangulation in $\mathbb{R}^3$, but we know no subdivision of space into a three-dimensional triangulation that conforms to a restricted CDT. (Note also that the homeomorphism that Edelsbrunner and Shah use is not the nearest point map; it is a different map based on the barycentric subdivision of the restricted Delaunay triangles. As the nearest point map is a natural way to connect a triangulation to the surface it represents—e.g., for the purpose of texture mapping—our result is interesting even just for restricted Delaunay triangulations.)

Our main new idea is a direct proof that, under the right sampling conditions, each restricted Voronoi cell is a "star-shaped" topological disk (Lemma 48 and Theorem 49). Another interesting part of this result is a proof that, under the right conditions, the nearest point map over any single restricted Delaunay triangle is an orientation-preserving homeomorphism from the triangle to its image on Σ (Theorem 57).

For the sake of brevity, we use the notation $|pq|$ to denote $d(p, q)$ throughout this section.

Restricted Voronoi Cells Are Topological Disks

Recall that for a site $v \in V$, v's *principal Voronoi cell* $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, v = \overline{\Sigma}_S \cap \mathrm{Vor}\,|_{\overline{\Sigma}} v$ excludes the portion of the extended restricted Voronoi cell on the extrusions. This section investigates sampling conditions that guarantee that each principal Voronoi cell has the topology of a closed disk. Corollary 50 shows that a constrained 0.44-sample suffices, whereas Theorem 49 gives a sampling condition more suitable for mesh generation algorithms. We begin with a simple circumstance in which an orthogonal projection is a homeomorphism.

Lemma 47. *Let Σ be a smooth surface without boundary in $\mathbb{R}^3$. Let $C \subset \mathbb{R}^3$ be a convex point set, let $D = C \cap \Sigma$, and suppose that D is connected. Let Γ be a plane in $\mathbb{R}^3$, and let φ be the continuous map that orthogonally projects $\mathbb{R}^3$ onto Γ. Suppose that for every point $u \in D$, the vector n_u normal to Σ at u is not parallel to Γ. Then the restricted projection $\varphi|_D$ is a homeomorphism from D to its image $\varphi(D)$ on Γ.*

Proof. First we show that $\varphi|_D$ is injective. Suppose for the sake of contradiction that two points $x, z \in D$ have $\varphi(x) = \varphi(z)$. All the points in D that map to $\varphi(x)$, including x and z, lie on a common line ℓ, perpendicular to Γ and passing through x. Let y be the point in $\Sigma \cap xz$ that is nearest x but is not x. (The point y might be z, or there might be a point closer to x.) As $D = C \cap \Sigma$ for a convex C, $y \in D$. No point in D is between x and y; that is, among the points in $D \cap \ell$, x and y are successive along ℓ.

As Σ is a 2-manifold without boundary in $\mathbb{R}^3$, it divides $\mathbb{R}^3$ into a bounded component and an unbounded, "outside" component. Let n_x and n_y be outward-facing vectors normal to Σ at x and y, respectively. Let n_Γ be a vector normal to Γ (and parallel to ℓ), directed so that $\angle(n_\Gamma, n_x) < 90°$. As we walk along ℓ, if x represents a transition from inside to outside, then y represents a transition from outside to inside (and vice versa), so $\angle(n_\Gamma, n_y) > 90°$.

As D is connected, there is a path $\gamma \subset D$ connecting x to y. As Σ is smooth, the outward-facing vector n_u normal to Σ at u varies continuously for $u \in \gamma$. Therefore, there is a point $v \in \gamma$ for which $\angle(n_\Gamma, n_v) = 90°$, which contradicts the assumption that for every $u \in D$, n_u is not parallel to Γ. Hence $\varphi|_D$ is injective.

Clearly φ is continuous. To see that the inverse $\varphi|_D^{-1}$ is continuous, consider the set L of all lines that are perpendicular to Γ and intersect a point in D. As Σ is smooth and no line in L intersects Σ tangentially at a point in D, the intersection $D \cap \ell$ varies continuously with $\ell \in L$. Therefore, $\varphi|_D^{-1}$ is continuous and $\varphi|_D$ is a homeomorphism from D to $\varphi(D)$. □

The following two results establish a "star-shaped" disk property that every principal Voronoi cell possesses in a constrained 0.44-sample. Consider a site $p \in V$ and its principal Voronoi cell $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$. Let F_λ be the open ball centered at p with radius $\lambda\,\mathrm{lfs}(p)$, and let $F = F_1$. By the definition of $\mathrm{lfs}(\cdot)$, F does not intersect the medial axis of Σ. By the Feature Ball Lemma (Lemma 39), $F_\lambda \cap \Sigma$ is a topological disk for every $\lambda < 1$.

Let n_p be a vector normal to Σ at p and let $T_p\Sigma$ be the plane tangent to Σ at p. Let φ be the continuous map that orthogonally projects $\mathbb{R}^3$ onto $T_p\Sigma$ (with the projection direction being parallel to n_p). By the Normal Variation Lemma (Lemma 21), for every point $x \in \Sigma$ with $|px| \le 0.9101\,\mathrm{lfs}(p)$, $\angle(n_p, n_x) < 90°$. Therefore by Lemma 47, for $\lambda \le 0.9101$, the orthogonal projection $\varphi|_{F_\lambda \cap \Sigma}$ is a homeomorphism from $F_\lambda \cap \Sigma$ to its image on $T_p\Sigma$.

Suppose that $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p \subset F_\lambda$ for some $\lambda \le 0.9101$. Let I_p be the image of $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$ under φ; so $\varphi|_{\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p}$ is a homeomorphism from $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$ to I_p. Note that $\varphi(p) = p$ (because p lies on $T_p\Sigma$) and p lies in the interior of I_p (where the "interior" is defined with respect to the tangent plane $T_p\Sigma$). The forthcoming Lemma 48 shows that for $\lambda \le 0.786151$, I_p is *star-shaped*: for every point $y \in I_p$, the line segment connecting y to p is a subset of I_p. Moreover, every point on yp except possibly y lies in the interior of I_p. The subsequent Theorem 49 uses Lemma 48 to show that I_p is homeomorphic to a disk, and thus so is $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$,

Consider a point $x \in \mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$, where $x \ne p$ and $x \in F_\lambda$ (hence $p, x \in F_\lambda \cap \Sigma$). Consider the line segment $p\varphi(x)$ on $T_p\Sigma$. Let $\gamma = \varphi|_{F_\lambda \cap \Sigma}^{-1}(p\varphi(x))$ be the unique curve on $F_\lambda \cap \Sigma$ that projects to the line segment $p\varphi(x)$; γ's endpoints are the site p and the point x.

Lemma 48. *Consider a site $p \in V$ and a point $x \in \mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$, where $x \ne p$. Suppose that $|px| < \xi\,\mathrm{lfs}(p)$, where $\xi = \sqrt{\frac{\sqrt{5}-1}{2}} \doteq 0.786151$. Define the orthogonal projection φ, the image $I_p = \varphi(\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p)$, and the curve $\gamma \subset \Sigma$ as described above. Then $\gamma \subset \mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$; moreover, every point in $\gamma \setminus \{x\}$ is in the interior of $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$ (where the "interior" is defined relative to Σ). Equivalently (by the homeomorphism φ), $p\varphi(x) \subset I_p$ and every point in $p\varphi(x) \setminus \{\varphi(x)\}$ is in the interior of I_p (where the "interior" is defined relative to $T_p\Sigma$).*

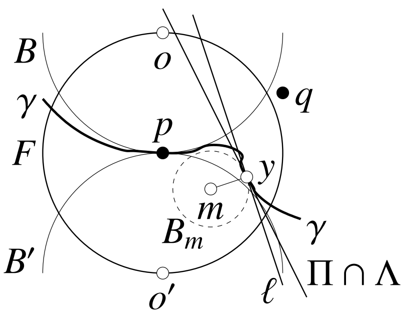

Figure A.1: In this figure, Π is the plane of the page and $p, x, y, o, o' \in \Pi$. The surface Σ and the curve $\gamma \subset \Sigma$ cannot intersect the open balls B, B', and B_m, so B and B_m are disjoint. Note that the center m of the medial ball B_m does not necessarily lie on Π (the page), and m cannot lie in the open ball F. The dashed circle is the boundary of $B_m \cap \Pi$. The line labeled $\Lambda \cap \Pi$ shows where pq's bisecting plane Λ intersects Π (the page), but Λ is not necessarily orthogonal to Π and q is not necessarily on Π. Neither $\Lambda \cap \Pi$ nor the line ℓ tangent to γ at y can separate o, o', and p from each other.

Furthermore, no point in $\gamma \setminus \{x\}$ is in any other site's restricted Voronoi cell, and for every other cell $\text{Vor}|_{\overline{\Sigma}_S} q$ *that contains x, γ is not tangent at x to the plane bisecting the line segment pq (hence Σ also is not tangent at x to the bisector).*

Proof. Suppose for the sake of contradiction that one of the following is true: $\gamma \not\subset \text{Vor}|_{\overline{\Sigma}_S} p$, or a point in $\gamma \setminus \{x\}$ is on the boundary of $\text{Vor}|_{\overline{\Sigma}_S} p$, or a point in $\gamma \setminus \{x\}$ is shared with another site's principal Voronoi cell, or for some site q such that $x \in \text{Vor}|_{\overline{\Sigma}_S} q$, γ is tangent at x to the plane bisecting the line segment pq.

In the first case ($\gamma \not\subset \text{Vor}|_{\overline{\Sigma}_S} p$), while we slide along γ from x to p, we encounter a point $y \in \gamma$ where γ leaves $\text{Vor}|_{\overline{\Sigma}_S} p$. By Lemma 44, the cell boundary is not caused by a portal shadow, so it must be caused by a bisector between sites. That is, there is a site $q \in V$ such that, letting Λ be the bisector of pq, γ crosses Λ at y (entering q's side of Λ), as illustrated in Figure A.1. If the first case does not apply ($\gamma \subset \text{Vor}|_{\overline{\Sigma}_S} p$) but one of the other three cases does apply, there is a point $y \in \gamma$ where γ intersects a bisector tangentially; that is, there is a $q \in V$ such that, letting Λ be the bisector of pq, γ intersects Λ tangentially at y. In all cases, let Λ_p be the closed halfspace that is bounded by Λ and contains p.

Let B and B' be the two open balls of radius $\text{lfs}(p)$ tangent to Σ at p, and let o and o' be their centers, respectively. Neither ball intersects Σ, hence neither ball contains any site. Therefore, p is a closest site to o and o', so $o \in \Lambda_p$ and $o' \in \Lambda_p$, as illustrated.

Let Π be the plane that contains p and x and is parallel to n_p (and thus orthogonal to $T_p\Sigma$) and observe that $I_p \subset \Pi$ and $\gamma \subset \Pi$. (Π is the plane of the page in Figure A.1.) The plane Π contains the points p, o, o', x, and y (but not necessarily q). Let $\Pi_p = \Pi \cap \Lambda_p$, a closed halfplane that also contains p, o, o', x, and y, with y on its boundary $\Pi \cap \Lambda$.

Let $\ell \subset \Pi$ be the line that is tangent to γ at y. Let the *vertical axis* be the line through o, p, and o'. We have seen that either γ intersects Λ tangentially at y, in which case $\ell = \Pi \cap \Lambda$, or γ leaves Λ_p at y, in which case ℓ is closer to being vertical (parallel to the vertical axis; vertical from the perspective of Figure A.1) than $\Pi \cap \Lambda$ is. More precisely, ℓ intersects the vertical axis farther from p than $\Pi \cap \Lambda$ does. Let ℓ_p be the closed halfplane with boundary ℓ that contains p; it follows that $o \in \ell_p$ and $o' \in \ell_p$, just as o, o', and p are all in Π_p.

There are two open medial balls that are tangent to Σ at y. Call them B_m and $B_{m'}$ and let m and m' be their centers. (Note that m and m' do not necessarily lie on Π.) As $\gamma \subset \Sigma$ and γ is a smooth curve, B_m and $B_{m'}$ are also tangent to γ and ℓ at y, and the line segment mm' is perpendicular to γ and ℓ at y. As y lies in the relative interior of mm', $\angle pym + \angle pym' = 180°$. Choose the labels m and m' so that $\angle pym < 90°$ and $\angle pym' > 90°$. (The inequalities are strict because $p \notin \ell$, so we can have $\angle pym = 90°$ only if ym is perpendicular to Π, but that is not possible as $\angle(n_p, n_y) < 90°$ for all $y \in \gamma \subset F_\xi$ by the Normal Variation Lemma (Lemma 21). Note that $B_{m'}$ can degenerate to an open halfspace, but B_m cannot because $\angle pym < 90°$ and $p \notin B_m$.)

Let Γ be the plane through y orthogonal to ym, and observe that $\ell \subset \Gamma$. As $\angle pym < 90°$, p and m are on the same side of Γ. As p and o are in the halfplane ℓ_p, either $o \in \Gamma$ or o is on the same side of Γ as p and m, and thus $\angle oym \leq 90°$. By Pythagoras' Theorem, $|oy|^2 + |my|^2 \geq |om|^2$ and $|py|^2 + |my|^2 > |pm|^2$. Observe that $m \notin F$, as m lies on the medial axis, which is disjoint from F. Hence $|pm| \geq \mathrm{lfs}(p)$ and thus $|py|^2 + |my|^2 > \mathrm{lfs}(p)^2$.

The surface Σ intersects none of the open balls B, B', or B_m, but it passes between B and B' at p. As Σ has no boundary and divides space into two pieces, one containing B and one containing B', the ball B_m must lie in one of those two pieces. Suppose without loss of generality that B_m lies in the same piece as B', as illustrated; then B_m must be disjoint from B. The radii of B and B_m are $\mathrm{lfs}(p)$ and $|my|$ respectively, so $|om| \geq \mathrm{lfs}(p) + |my|$. Combining this with the inequality $|oy|^2 + |my|^2 \geq |om|^2$ gives $|oy|^2 \geq \mathrm{lfs}(p)^2 + 2\,\mathrm{lfs}(p)\,|my|$. Combining this with the inequality $|my|^2 > \mathrm{lfs}(p)^2 - |py|^2$ gives $|oy|^2 > \mathrm{lfs}(p)^2 + 2\,\mathrm{lfs}(p)\sqrt{\mathrm{lfs}(p)^2 - |py|^2}$.

Create a coordinate system with p at the origin such that y_v is the vertical coordinate of y (parallel to the vertical axis; vertical in Figure A.1) and y_h is the horizontal coordinate of y (the axis in Π perpendicular to the vertical axis; horizontal in Figure A.1). Then $|oy|^2 + |o'y|^2 = y_h^2 + (y_v - \mathrm{lfs}(p))^2 + y_h^2 + (y_v + \mathrm{lfs}(p))^2 = 2y_h^2 + 2y_v^2 + 2\,\mathrm{lfs}(p)^2 = 2\,|py|^2 + 2\,\mathrm{lfs}(p)^2$. As $y \notin B'$, $|o'y|^2 \geq \mathrm{lfs}(p)^2$. Combining these with the inequality $|oy|^2 > \mathrm{lfs}(p)^2 + 2\,\mathrm{lfs}(p)\sqrt{\mathrm{lfs}(p)^2 - |py|^2}$ gives $|py|^2 = (|oy|^2 + |o'y|^2 - 2\,\mathrm{lfs}(p)^2)/2 > \mathrm{lfs}(p)\sqrt{\mathrm{lfs}(p)^2 - |py|^2}$. Recall that $y \in F_\xi$ because $y \in \gamma$; thus $|py| < \xi\,\mathrm{lfs}(p)$. It follows that $\xi^2 > \sqrt{1 - \xi^2}$, which is equivalent to $\xi^4 + \xi^2 - 1 > 0$, which implies that $\xi > \sqrt{\frac{\sqrt{5}-1}{2}}$. The result follows by contradiction. □

Theorem 49. *Let $p \in V$ be a site. Suppose that for every point $x \in \mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$, $|px| < \xi\,\mathrm{lfs}(p)$, where $\xi = \sqrt{(\sqrt{5} - 1)/2} \doteq 0.786151$. Then $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$ is homeomorphic to a closed disk. Moreover, no other restricted Voronoi cell intersects the interior of $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$.*

Proof. For every point $x \in \mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$, there is a path in $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$ connecting x to p by Lemma 48, so $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$ is connected. $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$ is the intersection of a smooth manifold Σ with a finite number of closed halfspaces, so $\mathrm{Vor}|_{\overline{\Sigma}_S}\, p$ is closed (both with respect to the manifold Σ and with respect to the ambient metric space $\mathbb{R}^3$). The site p does not lie on the boundary of $\mathrm{Vor}|_{\overline{\Sigma}_S}\, p$, as no bisector between p and another site can touch p.

Let $\bar{F}_\xi$ be the closed ball centered at p with radius ξ (whereas F_ξ is the open ball), and let I_F be the image of $\bar{F}_\xi \cap \Sigma$ under φ (the orthogonal projection onto $T_p\Sigma$). Recall that φ is injective over $\bar{F}_\xi \cap \Sigma$; we can define an inverse φ^{-1} over I_F, and the inverse is continuous. Recall I_p, the image of $\mathrm{Vor}|_{\overline{\Sigma}_S}\, p$ under φ. As $\mathrm{Vor}|_{\overline{\Sigma}_S}\, p \subseteq F_\xi \cap \Sigma$, it follows that $I_p \subset I_F$, φ is injective over $\mathrm{Vor}|_{\overline{\Sigma}_S}\, p$, and the inverse φ^{-1} is defined and continuous over I_p. As φ (restricted to $\mathrm{Vor}|_{\overline{\Sigma}_S}\, p$) is a homeomorphism from $\mathrm{Vor}|_{\overline{\Sigma}_S}\, p$ to I_p, we will show that the Voronoi cell is a topological disk by showing that I_p is a topological disk.

By continuity, I_p is closed and connected like the Voronoi cell, and it is bounded (hence compact) as well. By Lemma 48, I_p is *star-shaped* and for every point $y \in I_p$, every point on the segment yp except possibly y lies in the interior of I_p. As p lies in the interior of I_p, every radial line segment $pz \subset I_F$ connecting p to a point z on the boundary of I_F intersects one and only one point on the boundary of I_p. Thus we can define a function f that maps points on the boundary of I_F to points on the boundary of I_p. For convenience, we also define $f_\mathbb{R}$ to map each point z on the boundary of I_F to the distance $|pf(z)|$. If $f_\mathbb{R}$ is continuous, then I_p is homeomorphic to a unit disk (and we can construct an explicit homeomorphism by stretching or shrinking each line segment radiating out from p), and the theorem follows.

We now show that $f_\mathbb{R}$ is continuous. Let z be a point on the boundary of I_F; we show that $f_\mathbb{R}$ is continuous at z by showing that for every sufficiently small $\delta > 0$, there is an open arc a on the boundary of I_F with z in its relative interior such that every point $y \in a$ has $f_\mathbb{R}(y) \in (f_\mathbb{R}(z)-\delta, f_\mathbb{R}(z)+\delta)$. Let $x = f(z)$, which lies on the boundary of I_p at a distance of $f_\mathbb{R}(z)$ from p in the direction of z. Let x^- be the point at a distance of $f_\mathbb{R}(z) - \delta$ from p in the direction of z, and let x^+ be the point at a distance of $f_\mathbb{R}(z) + \delta$ from p in the direction of z. (We must take δ sufficiently small that $f_\mathbb{R}(z) - \delta > 0$ and $f_\mathbb{R}(z) + \delta$ is small enough that $x^+ \in I_f$.) By Lemma 48, x^- is in the interior of I_p. As I_p is star-shaped with x on its boundary, $x^+ \notin I_p$. As I_p is closed, the circle of radius $f_\mathbb{R}(z) + \delta$ centered at p includes an open arc a^+ that has x^+ in its relative interior and does not intersect I_p. As x^- is in the interior, the circle of radius $f_\mathbb{R}(z) - \delta$ centered at p includes an open arc a^- that is entirely in the interior of I_p and has x^- in its relative interior. By comparing the angle intervals of a^- and a^+ and taking their intersection (the common angles), we produce a curve a on the boundary of I_F with z in its relative interior such that every point $y \in a$ has $f(y)$ trapped between a^+ and a^-, hence $f_\mathbb{R}(y) \in (f(z) - \delta, f(z) + \delta)$. Therefore $f_\mathbb{R}$ is continuous, f is continuous, and $\mathrm{Vor}|_{\overline{\Sigma}_S}\, p$ is a topological closed disk.

To prove the second claim—that no other restricted Voronoi cell intersects the interior of $\mathrm{Vor}|_{\overline{\Sigma}_S}\, p$—observe that for any point y in the interior, $\varphi(y)$ lies in the interior of I_p, and we can extend the line segment $p\varphi(y)$ to a point $\varphi(x)$ on the boundary of I_p. Define $\gamma = \varphi|^{-1}_{F_\lambda \cap \Sigma}(p\varphi(x))$. By Lemma 48, no point in $\gamma \setminus \{x\}$ is in the restricted Voronoi cell of any site besides p. Hence y is not in any other restricted Voronoi cell. □

It is notable that the condition $|px| < \xi\, \mathrm{lfs}(p)$ of Lemma 48 and Theorem 49 implies, by the Normal Variation Lemma (Lemma 21), that $\angle(n_p, n_x) < \eta(\xi) = 60°$. The sixty-degree bound is exact. (We do not have any intuition for why that number comes out so cleanly.)

We now determine which constrained ϵ-samples satisfy the condition.

Corollary 50. *Let V be an constrained ϵ-sample of (Σ, S, Z) for $\epsilon < \frac{\xi}{\xi+1} \doteq 0.440137$. Every principal Voronoi cell in $\mathrm{Vor}\,|_{\bar{\Sigma}_S} V$ is homeomorphic to a closed disk. Moreover, no restricted Voronoi cell intersects the interior of another restricted Voronoi cell. Moreover, there is no point $y \in \mathrm{Vor}\,|_{\bar{\Sigma}_S}\, p \cap \mathrm{Vor}\,|_{\bar{\Sigma}_S}\, q$ at which Σ is tangent to the bisector between two sites p and q.*

Proof. Consider any site $p \in V$ and any point $x \in \mathrm{Vor}\,|_{\bar{\Sigma}_S}\, p$. As V is a constrained ϵ-sample, $|px| \leq \epsilon\, \mathrm{lfs}(x)$. By the Feature Translation Lemma (Lemma 38), $|px| \leq (\epsilon/(1-\epsilon))\, \mathrm{lfs}(p) < \xi\, \mathrm{lfs}(p)$. The first two claims follow by Theorem 49. The third claim follows from Lemma 48. □

The following corollary shows that if we extend the condition of Theorem 49 to all principal Voronoi cells (or impose the condition of Corollary 50), every connected component of Σ has at least six sites on it.

Corollary 51. *Let $V \subset \Sigma$ be a nonempty, finite set of points (sites) on Σ. Suppose that for every site $p \in V$ and every point $x \in \mathrm{Vor}\,|_{\bar{\Sigma}_S}\, p$, $|px| < \xi\, \mathrm{lfs}(p)$, where $\xi = \sqrt{(\sqrt{5}-1)/2} \doteq 0.786151$. (Alternatively, suppose that V is a constrained ϵ-sample of (Σ, S, Z) for $\epsilon < \frac{\xi}{\xi+1} \doteq 0.440137$.)*

Then every connected component of Σ has at least six sites and at least six principal Voronoi cells on it.

Proof. By Theorem 49 (or Corollary 50), every principal Voronoi cell in $\mathrm{Vor}\,|_{\bar{\Sigma}_S} V$ is homeomorphic to a closed disk. Therefore, every principal Voronoi cell lies on just one connected component of Σ. By the Normal Variation Lemma (Lemma 21), for every site $p \in V$ and every point $x \in \mathrm{Vor}\,|_{\bar{\Sigma}_S}\, p$, $\angle(n_p, n_x) < \eta(\xi) = 60°$.

Let $\mathring{\Sigma}$ be a connected component of Σ. For any unit vector u on the unit sphere, let y be a point on $\mathring{\Sigma}$ that is most extreme in the direction u. As $\mathring{\Sigma}$ is a smooth surface without boundary, u is normal to $\mathring{\Sigma}$ at y and oriented to the outside of $\mathring{\Sigma}$; that is, the unit vector $n_y = u$ is an outward-facing normal vector at y. It follows that the outside-facing unit normal vectors on $\mathring{\Sigma}$ constitute the entire sphere of directions. A principal Voronoi cell $\mathrm{Vor}\,|_{\bar{\Sigma}_S}\, p$ can contain only points on Σ whose outside-facing normals are less than 60° from n_p. At least six sites are required on a sphere so that every point on the sphere is less than 60° from one of the sites; five do not suffice [44]. (The six points where the coordinate axes intersect the unit sphere suffice.) Hence there are at least six sites and six principal Voronoi cells on $\mathring{\Sigma}$.

The same reasoning applies to every connected component of Σ. □

The Nearest Point Map and Circumscribing Spheres

Recall that the nearest point map ν maps any point $x \in \mathbb{R}^3 \setminus M$ to the point $\tilde{x} = \nu(x)$ nearest x on Σ. The following lemma helps to constrain where $\tilde{x}$ can lie. We will use it later for two purposes: to prove conditions under which a mapped triangle $\nu(\tau)$ does not intersect any site other than τ's

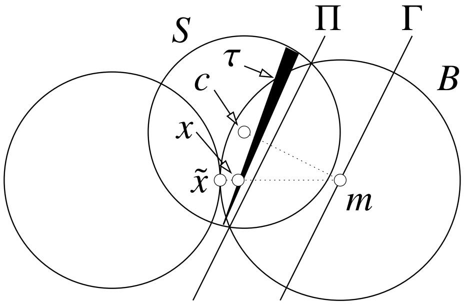

Figure A.2: For every point $x \in \tau$ except τ's vertices, $\tilde{x}$ is strictly inside S.

vertices, and to help prove conditions under which ν defines a homeomorphism from τ to its image $\nu(\tau)$ on Σ.

Lemma 52. *Consider three non-collinear points $p, p', p'' \in \Sigma$ and the triangle $\tau = \triangle pp'p''$. Let S be a sphere that passes through p, p', and p''. Let r be the radius of S, and suppose that $r \leq \mathrm{lfs}(p)/2$. Then for every point $x \in \tau \setminus \{p, p', p''\}$, $\tilde{x} = \nu(x)$ is strictly inside S.*

Proof. Consider a point $x \in \tau \setminus \{p, p', p''\}$. As τ's vertices lies on S, x is inside S. If $\tilde{x} = x$ the lemma follows immediately, so suppose that $\tilde{x} \neq x$ and thus $x \notin \Sigma$. There are two open medial balls tangent to Σ at $\tilde{x}$; let B be the one that contains x, as illustrated in Figure A.2. Let m be the center of B; m lies on the medial axis of Σ. Observe that x lies on the line segment $m\tilde{x}$.

If the entire closure of B is strictly inside S, the lemma follows immediately; so assume it is not. The entirety of B cannot be outside S, as $x \in B$ and x is inside S. Nor is $S \subset B$ possible, as τ's vertices are not in B. Hence the intersection of S with B's boundary is a circle or a point. If it is a circle, let Π be the affine hull of that circle, as illustrated; if it is a point, let Π be the plane tangent to S and B at that point. Let $\bar{\Pi}_S$ be the closed halfspace bounded by Π that includes $S \setminus B$, and let Π_S be the open version of the same halfspace. The portion of B's boundary in Π_S is entirely enclosed by S. The portion of S in the open halfspace complementary to $\bar{\Pi}_S$ is entirely included in B. Every vertex of τ lies on S but not in B, hence τ's vertices lie in $\bar{\Pi}_S$. Therefore, $\tau \subset \bar{\Pi}_S$ and $x \in \bar{\Pi}_S$.

Let c be the center of S. Observe that the plane Π is orthogonal to the line segment cm. By assumption, $r \leq \mathrm{lfs}(p)/2$, so $|pm| \geq \mathrm{lfs}(p) \geq 2r$. As p lies on S and $|pm|$ is at least twice the radius r of S, it follows that m lies outside of S or on S. Hence $m \notin \bar{\Pi}_S$.

Given the facts that x lies on the line segment $m\tilde{x}$, $m \notin \bar{\Pi}_S$, $x \in \bar{\Pi}_S$, and $\tilde{x} \neq x$, it follows that $\tilde{x} \in \Pi_S$. As $\tilde{x}$ is also on B's boundary, $\tilde{x}$ is strictly inside S. □

A corollary of Lemma 52 is that if the vertices of a restricted Delaunay triangle τ are sufficiently close to τ's dual restricted Voronoi vertex, then the image of τ under the nearest point map ν does not intersect any site other than τ's vertices.

Corollary 53. *Let V be a finite set of points on Σ. Let $p, p', p'' \in V$ be three sites that generate a restricted Voronoi vertex $u \in \mathrm{Vor}|_\Sigma V$ and its dual restricted Delaunay triangle $\tau = \triangle pp'p''$. Suppose that $|pu| \leq \mathrm{lfs}(p)/2$. (Note that $|pu| = |p'u| = |p''u|$.) Then $\nu(\tau)$ intersects no site in $V \setminus \{p, p', p''\}$.*

Proof. Suppose for the sake of contradiction that for some site $w \in V \setminus \{p, p', p''\}$, $w \in \nu(\tau)$. Let $x \in \tau$ be the point for which $\nu(x) = w$. As τ is a restricted Delaunay triangle dual to u, τ's vertices lie on a sphere S that has center u and encloses no site, particularly not w. The radius of S is $|pu|$. By Lemma 52, w is strictly inside S. The result follows by contradiction. □

The Nearest Point Map on a Triangle is a Homeomorphism

The forthcoming Theorem 57 establishes conditions under which the nearest point map, restricted to a restricted Delaunay triangle, is a homeomorphism; so there are no foldovers within a single triangle's map.

Assuming Σ is a 2-manifold, we define a *restricted Voronoi vertex* to be a point in the intersection of three distinct extended Voronoi cells. However, without suitable sampling conditions, such an intersection might include one or more line segments. Theorem 57 also establishes conditions that guarantee that a restricted Voronoi vertex is isolated from any other points in the intersection, thereby justifying the name "vertex." We start with a technical lemma.

Lemma 54. *Let u be a point on Σ. Let $\tau \subset \mathbb{R}^3$ be a simplex. Let $x \in \tau$ be a point that does not lie on the medial axis of Σ. Let $\tilde{x}$ be the point on Σ nearest x. There is a vertex p of τ such that $|p\tilde{x}| \leq |pu|$, and such that $|p\tilde{x}| < |pu|$ if $\tilde{x} \neq u$.*

Proof. If $\tilde{x} = u$ then the result follows immediately, so assume that $\tilde{x} \neq u$. As x does not lie on the medial axis, $\tilde{x}$ is the unique point on Σ nearest x. As u also lies on Σ, $|x\tilde{x}| < |xu|$. Let Π be the plane that bisects the line segment $\tilde{x}u$, and observe that x lies on the same side of Π as $\tilde{x}$. As $x \in \tau$ and τ is a simplex, some vertex p of τ lies on the same side of Π as $\tilde{x}$, thus $|p\tilde{x}| < |pu|$. □

Given an extended restricted Voronoi vertex u and its dual restricted Delaunay triangle τ, the following two lemmas relate the normal vectors n_u and $n_{\tilde{x}}$ at any point $\tilde{x} \in \nu(\tau)$, showing that all these normals point to the same side of τ. Recall that a *principal vertex* is an extended restricted Voronoi vertex that lies on $\overline{\Sigma}_S$. A *secondary vertex* is an extended restricted Voronoi vertex that is not principal; it lies on an extrusion but not on a portal curve.

Lemma 55. *Let u be a principal vertex and let $\tau = \triangle pp'p''$ be its dual restricted Delaunay triangle. Let x be any point on τ, and let $\tilde{x} = \nu(x)$ be the point on Σ nearest x. Let n_u be an outward-facing vector normal to Σ at u, let $n_{\tilde{x}}$ be an outward-facing vector normal to Σ at $\tilde{x}$, and let n_τ be a vector normal to τ. Let $R = |pu| = |p'u| = |p''u|$, and suppose that $R \leq 0.3202\,\mathrm{lfs}(u)$.*

Then the angles $\angle(n_u, n_\tau)$ and $\angle(n_{\tilde{x}}, n_\tau)$ are either both less than 90° or both greater than 90° (depending on which way n_τ is directed). Equivalently, the dot products $n_u \cdot n_\tau$ and $n_{\tilde{x}} \cdot n_\tau$ are either both positive or both negative.

Proof. Let r be τ's circumradius. Let S be the sphere with center u and radius R, which passes through all three vertices of τ. As τ's circumcircle is a cross section of S, $r \leq R$.

Suppose without loss of generality that p is the vertex of τ nearest $\tilde{x}$. Let $q \in \{p, p', p''\}$ be the vertex at τ's largest plane angle. As $|pu| = |qu| = R \leq 0.3202\,\mathrm{lfs}(u)$, by the Feature Translation Lemma (Lemma 38), $\mathrm{lfs}(u) \leq \mathrm{lfs}(p)/(1 - 0.3202)$ and likewise $\mathrm{lfs}(u) \leq \mathrm{lfs}(q)/(1 - 0.3202)$, so

$r \le R \le \frac{0.3202}{0.6798}\,\mathrm{lfs}(p) < 0.48\,\mathrm{lfs}(p)$ and likewise $r \le \frac{0.3202}{0.6798}\,\mathrm{lfs}(q)$. By Lemma 52 (with S as defined above), $|\tilde{x}u| \le R$; hence $|\tilde{x}u| \le 0.3202\,\mathrm{lfs}(u)$. By Lemma 54, there is a vertex $\dot{p}$ of τ such that $|\dot{p}\tilde{x}| \le |\dot{p}u|$; hence $|p\tilde{x}| \le |\dot{p}\tilde{x}| \le |\dot{p}u| = R \le \frac{0.3202}{0.6798}\,\mathrm{lfs}(p)$. By the Normal Variation Lemma (Lemma 21), $\angle(n_p, n_u) \le \eta(0.3202)$, $\angle(n_q, n_u) \le \eta(0.3202)$, $\angle(n_{\tilde{x}}, n_u) \le \eta(0.3202)$, and $\angle(n_p, n_{\tilde{x}}) \le \eta(0.3202/0.6798)$, where $\eta(\delta) = \arccos\left(1 - \frac{\delta^2}{2\sqrt{1-\delta^2}}\right)$, $\eta(0.3202) < 18.94°$, and $\eta(0.3202/0.6798) < 29.05°$.

If τ's plane angle at the vertex p is 56.653° or greater, then by the Triangle Normal Lemma (Lemma 1), $\sin \angle(n_p, n_\tau) \le r \cot 28.3265°/\mathrm{lfs}(p) < (0.3202/0.6798) \cdot 1.8552 < 0.8739$. Therefore, either $\angle(n_p, n_\tau) < 60.92°$ or $\angle(n_p, n_\tau) > 119.08°$, depending on which way n_τ is directed. Suppose without loss of generality that n_τ is directed so that $\angle(n_p, n_\tau) < 60.92°$. Then $\angle(n_u, n_\tau) \le \angle(n_p, n_u) + \angle(n_p, n_\tau) < 18.94° + 60.92° = 79.86°$ and $\angle(n_{\tilde{x}}, n_\tau) \le \angle(n_p, n_{\tilde{x}}) + \angle(n_p, n_\tau) < 29.05° + 60.92° = 89.97°$, so both angles are less than 90° as claimed.

Otherwise, τ's plane angle at p is less than 56.653°, so τ's plane angle at q (τ's largest plane angle) is greater than $(180° - 56.653°)/2 = 61.6735°$. By the Triangle Normal Lemma (Lemma 1), $\sin \angle(n_q, n_\tau) \le r \cot 30.83675°/\mathrm{lfs}(q) < (0.3202/0.6798) \cdot 1.6751 < 0.7891$. Therefore, either $\angle(n_q, n_\tau) < 52.11°$ or $\angle(n_q, n_\tau) > 127.89°$, depending on which way n_τ is directed. Suppose without loss of generality that n_τ is directed so that $\angle(n_q, n_\tau) < 52.11°$. Then $\angle(n_u, n_\tau) \le \angle(n_q, n_u) + \angle(n_q, n_\tau) < 18.94° + 52.11° = 71.05°$ and $\angle(n_{\tilde{x}}, n_\tau) \le \angle(n_{\tilde{x}}, n_u) + \angle(n_q, n_u) + \angle(n_q, n_\tau) < 18.94° + 18.94° + 52.11° = 89.99°$, confirming that both angles are less than 90°. □

The following lemma applies to all extended restricted Voronoi vertices, both principal and secondary. Although a secondary vertex u does not lie on Σ, but rather on an extrusion, we still speak of an "outward-facing" normal vector n_u consistent with the outward-facing normal vectors on Σ, as we can extend Σ's orientation onto the extrusions.

Lemma 56. *Let u be an extended Voronoi vertex (principal or secondary) and let $\tau = \triangle pp'p''$ be the restricted Delaunay triangle dual to u, where p is the vertex of τ at τ's largest plane angle. Let μ be the positive root of $4\mu^4 = (1 - 4\mu^2)(1 - \sqrt{3}\mu)^2$, with approximate value $\mu \doteq 0.3606001$. Let $R = |pu| = |p'u| = |p''u|$ and suppose that $R < \mu\,\mathrm{lfs}(p)$. If u is a secondary vertex on an extrusion of a segment s, suppose also that the length of s is at most $\rho\,\mathrm{lfs}(a)$, where $\rho \le 0.47$ and a is an endpoint of s. Let x be any point on τ, and let $\tilde{x} = \nu(x)$ be the point on Σ nearest x. Let n_u be an outward-facing vector normal to $\widetilde{\Sigma}$ at u, let $n_{\tilde{x}}$ be an outward-facing vector normal to Σ at $\tilde{x}$, and let n_τ be a vector normal to τ.*

Then the angles $\angle(n_u, n_\tau)$ and $\angle(n_{\tilde{x}}, n_\tau)$ are either both less than 90° or both greater than 90° (depending on which way n_τ is directed). Equivalently, the dot products $n_u \cdot n_\tau$ and $n_{\tilde{x}} \cdot n_\tau$ are either both positive or both negative.

Proof. Let r be τ's circumradius. Let S be the sphere with center u and radius R, which passes through all three vertices of τ. As τ's circumcircle is a cross section of S, $r \le R$.

By the Triangle Normal Lemma (Lemma 1), $\sin \angle(n_p, n_\tau) \le \sqrt{3}r/\mathrm{lfs}(p) < \sqrt{3}\mu$. Suppose without loss of generality that n_τ is directed so that $\angle(n_p, n_\tau)$ is acute; then $\angle(n_p, n_\tau) < 38.652°$.

By Lemma 52 (with S as defined above), $|u\tilde{x}| \le R$; hence $|p\tilde{x}| \le |pu| + |u\tilde{x}| \le 2R < 2\mu\,\mathrm{lfs}(p)$. By the Normal Variation Lemma (Lemma 21), $\angle(n_p, n_{\tilde{x}}) < \eta(2\mu)$ where $\eta(\delta) = \arccos\left(1 - \frac{\delta^2}{2\sqrt{1-\delta^2}}\right)$.

Therefore, $\angle(n_p, n_{\tilde{x}}) < \arccos\left(1 - \frac{2\mu^2}{\sqrt{1-4\mu^2}}\right) = \arccos\left(1 - (1 - \sqrt{3}\mu)\right) = \arccos(\sqrt{3}\mu)$ and $\angle(n_{\tilde{x}}, n_\tau) \leq \angle(n_p, n_{\tilde{x}}) + \angle(n_p, n_\tau) < \arccos(\sqrt{3}\mu) + \arcsin(\sqrt{3}\mu) = 90°$.

If u is a principal vertex, then as $|pu| = R < \mu \operatorname{lfs}(p)$, by the Normal Variation Lemma $\angle(n_p, n_u) < \eta(\mu) < 21.52°$. Therefore, $\angle(n_u, n_\tau) \leq \angle(n_p, n_u) + \angle(n_p, n_\tau) < 21.52° + 38.652° = 60.172°$. So $\angle(n_u, n_\tau)$ and $\angle(n_{\tilde{x}}, n_\tau)$ are both less than 90°, and the lemma holds.

If u is a secondary vertex, let s be the segment on whose extrusion u lies, let h_s be the cutting plane for s, and let $\zeta_s \subset h_s \cap \Sigma$ be the portal curve for s. Let $\bar{u} \in \zeta_s$ be the point nearest u on h_s, and note that $\bar{u} \in \zeta_s$ and $\bar{u} \in \Sigma$. It follow from Theorem 18 that the plane h_s separates τ from u; therefore, $|p\bar{u}| < |pu|$.

As $|p\bar{u}| < |pu| = R < \mu \operatorname{lfs}(p)$, by the Normal Variation Lemma $\angle(n_p, n_{\bar{u}}) < \eta(\mu) < 21.52°$. As the length of s is at most $\rho \operatorname{lfs}(a)$, we have $|a\bar{u}| \leq \rho \operatorname{lfs}(a)$ and, by the Normal Variation Lemma, $\angle(n_a, n_{\bar{u}}) \leq \eta(\rho) \leq \eta(0.47) < 28.971°$. As the site u lies on an extrusion from ζ_s, the vector n_u normal to the extrusion at u is the projection of $n_{\bar{u}}$ onto h_s. As h_s is parallel to n_a, $\angle(n_u, n_{\bar{u}}) \leq \angle(n_a, n_{\bar{u}}) < 28.971°$.

Therefore, $\angle(n_u, n_\tau) \leq \angle(n_u, n_{\bar{u}}) + \angle(n_p, n_{\bar{u}}) + \angle(n_p, n_\tau) < 28.971° + 21.52° + 38.652° = 89.143°$. Hence, $\angle(n_u, n_\tau)$ and $\angle(n_{\tilde{x}}, n_\tau)$ are both less than 90° as claimed. □

Theorem 57. *Consider* $\operatorname{Vor}|_{\tilde{\Sigma}} V$*, where* $\Sigma \subset \mathbb{R}^3$ *is a smooth* 2*-manifold without boundary,* $V \subset \Sigma$ *is a finite sample, and each segment* $s \in S$ *has length at most* $0.47 \operatorname{lfs}(a)$ *for some endpoint a of s. Consider three distinct sites* $p, p', p'' \in V$ *and the triangle* $\tau = \triangle pp'p''$*, where p is the vertex of* τ *at* τ*'s largest plane angle. Let* $U = \operatorname{Vor}|_{\tilde{\Sigma}} p \cap \operatorname{Vor}|_{\tilde{\Sigma}} p' \cap \operatorname{Vor}|_{\tilde{\Sigma}} p''$ *(which is the extended Voronoi face dual to* τ*) and suppose that* $U \neq \emptyset$ *(so* τ *is a restricted Delaunay triangle). Let u be a point in U and let* $R = |pu| = |p'u| = |p''u|$*. Suppose that at least one of the following holds: either u is a principal vertex and* $R \leq 0.3202 \operatorname{lfs}(u)$*, or* $R \leq 0.3606 \operatorname{lfs}(p)$*.*

Then the nearest point map ν *restricted to* τ*, denoted* $\nu|_\tau$*, is a homeomorphism from* τ *to its image* $\nu(\tau)$ *on* Σ*. Moreover, no normal segment of* Σ *that intersects* τ *is parallel to* τ*. Moreover, u is isolated from the other points in U.*

Proof. First we show that τ does not intersect the medial axis M of Σ, so the (restricted) nearest point map $\nu|_\tau$ is defined and continuous over τ. If the condition $R \leq 0.3202 \operatorname{lfs}(u)$ holds, then the distance from u to any point in τ is at most $0.3202 \operatorname{lfs}(u)$, whereas the distance from u to M is $\operatorname{lfs}(u)$, so τ is disjoint from M. If the condition $R \leq 0.3606 \operatorname{lfs}(p)$ holds, then the distance from p to any point in τ is at most $2R \leq 0.7212 \operatorname{lfs}(p)$, whereas the distance from p to M is $\operatorname{lfs}(p)$; again τ is disjoint from M. In either case, $\nu|_\tau$ is continuous over τ.

By Lemma 55 (if $R \leq 0.3202 \operatorname{lfs}(u)$) or Lemma 56 (if $R \leq 0.3606 \operatorname{lfs}(p)$), for every point $x \in \tau$, $\angle(n_{\tilde{x}}, n_\tau) \neq 90°$; therefore, the unique normal segment $\ell_{\tilde{x}}$ that passes through x is not parallel to τ. This proves our claim that no normal segment that intersects τ is parallel to τ. It follows that the nearest point map $\nu|_\tau$ is injective: if two distinct points $x, x' \in \tau$ could map to the same point $\tilde{x} \in \Sigma$, then $\tilde{x}$'s normal segment $\ell_{\tilde{x}}$ would intersect both x and x' and thus be parallel to τ, but that is not possible.

The nearest point map $\nu|_\tau$ is a continuous bijection between a compact set τ and its image $\nu(\tau)$ on a bounded manifold. Its inverse ν^{-1} is also continuous over $\nu(\tau)$, as the normal lines are a continuous function of the points on Σ, and the intersection of a line with τ's affine hull Π is a continuous function over the domain of lines that are not parallel to Π. Hence $\nu|_\tau$ is a homeomorphism (proving our first claim).

To address our third claim, let ℓ_τ be the line composed of all the points in the augmented three-dimensional space $\widetilde{X}$ that are equidistant to the sites p, p', and p'', and observe that $U \subset \ell_\tau$. By Lemma 55 or Lemma 56, $\angle(n_u, n_\tau) \neq 90°$. As ℓ_τ is parallel to the normal vector n_τ, ℓ_τ does not intersect Σ tangentially at u, so u is isolated from the other points in U. □

Extended Voronoi Edges Are Topological Line Segments

Theorem 49 and Corollary 50 give conditions under which the principal Voronoi cells are topological closed disks, and no cell intersects the interior of another cell. Theorem 57 gives conditions under which the intersection of any three distinct extended Voronoi cells is composed of isolated points, which we call "extended Voronoi vertices." What about an intersection of two distinct extended Voronoi cells? We call such an intersection an *extended Voronoi edge* if it contains a connected curve (and thus it is not merely a set of isolated points). The following lemma helps to justify this name.

Lemma 58. *Consider an extended Voronoi diagram* $\mathrm{Vor}|_{\widetilde{\widetilde{\Sigma}}}V$. *Suppose that every extended Voronoi cell is a topological closed disk and every intersection of three distinct extended Voronoi cells is a set of isolated points (i.e., no two distinct points are path-connected). Suppose also that no extended Voronoi cell intersects the interior of another extended Voronoi cell. Then for every pair of distinct sites* $p, q \in V$, $\mathrm{Vor}|_{\widetilde{\widetilde{\Sigma}}}p \cap \mathrm{Vor}|_{\widetilde{\widetilde{\Sigma}}}q$ *is either a topological circle containing no extended Voronoi vertex or a union of disjoint topological closed* 1*-balls and isolated points, where each isolated point is an extended Voronoi vertex and each* 1*-ball contains exactly two extended Voronoi vertices which are its endpoints.*

Moreover, if at least three sites in V lie on each connected component of Σ, *then the possibility that* $\mathrm{Vor}|_{\widetilde{\widetilde{\Sigma}}}p \cap \mathrm{Vor}|_{\widetilde{\widetilde{\Sigma}}}q$ *is a topological circle is eliminated, every extended Voronoi cell has at least two extended Voronoi vertices on its boundary, and every connected component of* Σ *has at least two extended Voronoi vertices on it.*

Proof. As $\widetilde{\widetilde{\Sigma}}$ is a surface without boundary and $\widetilde{\widetilde{\Sigma}}$ is also a union of extended Voronoi cells, which are topological closed disks, each point on the boundary of each extended Voronoi cell is shared with at least one other extended Voronoi cell. By assumption, no interior point of a cell is shared with another cell.

Consider a site p, its extended Voronoi cell $\mathrm{Vor}|_{\widetilde{\widetilde{\Sigma}}}p$, and the cell's boundary C, which is a topological circle. If two or more extended Voronoi vertices lie on C, they subdivide C into two or more topological closed 1-balls (as extended Voronoi vertices are isolated points). Let I be one of these topological 1-balls, or let $I = C$ if C contains fewer than two extended Voronoi vertices. The subset of I obtained by removing its extended Voronoi vertices is path-connected, so the points in that subset are all shared with one and only one other site q; hence $I \subseteq \mathrm{Vor}|_{\widetilde{\widetilde{\Sigma}}}p \cap \mathrm{Vor}|_{\widetilde{\widetilde{\Sigma}}}q$. It follows that for every $p, q \in V$, $\mathrm{Vor}|_{\widetilde{\widetilde{\Sigma}}}p \cap \mathrm{Vor}|_{\widetilde{\widetilde{\Sigma}}}q$ is either a topological circle or a union of 1-balls and extended Voronoi vertices. The intersection of two cells cannot include two 1-balls that are not disjoint—that is, two 1-balls that share a extended Voronoi vertex—because the shared vertex lies on the boundary of a third cell, which implies that at least one of the three cells is not a topological closed disk.

For the same reason, if the intersection of two cells is a topological circle, no extended Voronoi vertex can lie on the circle. This establishes the lemma's first claim.

If the intersection of two cells is a topological circle (with no extended Voronoi vertex), then as the two cells are topological disks, their union is a topological sphere covering an entire connected component of Σ. Hence, if at least three sites in V lie on each connected component of Σ, then no two cells have a circle as their intersection. The lemma's second claim follows. □

The Nearest Point Map Is Surjective

Theorem 57 and the forthcoming Lemmas 59, 60, and 62 use the idea of assigning an orientation to Σ, which manifests as both a normal vector direction and a rotary spin in the tangent space. There are two directions in which a normal vector n_u can point; let us choose the normal vectors so they all point to the same side of Σ. (Note: we don't actually need Σ to be orientable; it suffices to examine one patch that is orientable in isolation, such as a restricted Voronoi cell $\operatorname{Vor}|_\Sigma w$.) Then we define a counterclockwise ordering of the restricted Voronoi cells adjoining a restricted Voronoi vertex u according to a *right-hand rule*: with the thumb of your right hand pointing in the direction of n_u, your fingers curl in a direction that defines the *counterclockwise* ordering of cells adjoining u.

We can extend this notion of orientation to all the points on the normal segments. Each point $y \in \mathbb{R}^3 \setminus M$ lies on the normal segment ℓ_x of a point $x = \nu(y) \in \Sigma$. The point y inherits both aspects of x's orientation: the orientation direction n_x, parallel to ℓ_x, and the counterclockwise ordering around ℓ_x, derived from n_x by the right-hand rule.

Let $\tau \subset \mathbb{R}^3 \setminus M$ be a triangle whose vertices lie on Σ. Theorem 57 guarantees (under the stated conditions) that the nearest point map ν is a homeomorphism from τ to $\nu(\tau)$. By Lemma 55 and 56, the orientations of the points on τ are all mutually consistent; their orientation directions all point to the same side of τ. Let n_τ be a unit vector orthogonal to τ that points to the same side as well. The right-hand rule induces a counterclockwise ordering of τ's vertices around τ's boundary, which is also a counterclockwise ordering of $\nu(\tau)$'s vertices around $\nu(\tau)$'s boundary.

Consider a restricted Voronoi vertex $u \in \operatorname{Vor}|_\Sigma V$ generated by three sites $p, p', p'' \in V$ and its dual restricted Delaunay triangle $\tau = \triangle pp'p''$. Let ℓ_τ be the line comprising the points equidistant from p, p', and p''; ℓ_τ passes through both τ's circumcenter and u. (However, ℓ_τ is *not* necessarily parallel to n_u.) Imagine the three-site Voronoi diagram $\operatorname{Vor}\{p, p', p''\}$: it subdivides $\mathbb{R}^3$ into three wedges W_p, $W_{p'}$, and $W_{p''}$ whose mutual intersection is ℓ_τ. The restricted Voronoi cells in $\operatorname{Vor}|_\Sigma V$ satisfy the inclusions $\operatorname{Vor}|_\Sigma p \subset W_p$, $\operatorname{Vor}|_\Sigma p' \subset W_{p'}$, and $\operatorname{Vor}|_\Sigma p'' \subset W_{p''}$. Therefore, the cyclical ordering of τ's vertices around ℓ_τ is consistent with the cyclical ordering of their restricted Voronoi cells around ℓ_τ where they touch u. You might expect that if the former is counterclockwise, then so is the latter. However, if the surface is twisted enough that the angle between n_u and n_τ exceeds 90°, this expectation is violated: if τ's vertices p, p', and p'' occur in counterclockwise order around τ's perimeter, then $\operatorname{Vor}|_\Sigma p$, $\operatorname{Vor}|_\Sigma p'$, $\operatorname{Vor}|_\Sigma p''$ are in clockwise order around u. Intuitively, this causes a nasty “foldover” in the restricted Delaunay triangulation, which may prevent ν from being injective over the triangulation. We shall show that such foldovers can be prevented by use of a sufficiently dense sample.

Lemma 59. *Let u be a principal vertex and let $\tau = \triangle pp'p''$ be its dual restricted Delaunay triangle. Let $R = |pu| = |p'u| = |p''u|$ and suppose that at least one of the following holds: either $R \leq 0.3202\,\mathrm{lfs}(u)$ or $R \leq 0.3606\,\mathrm{lfs}(q)$ for each $q \in \{p, p', p''\}$. Then p, p', and p'' are in counterclockwise order around τ if and only if $\mathrm{Vor}|_{\tilde{\Sigma}}p$, $\mathrm{Vor}|_{\tilde{\Sigma}}p'$, and $\mathrm{Vor}|_{\tilde{\Sigma}}p''$ adjoin u in counterclockwise order around u.*

Proof. By Lemma 55 (if $R \leq 0.3202\,\mathrm{lfs}(u)$) or Lemma 56 (if $R \leq 0.3606\,\mathrm{lfs}(q)$ for each q), the outward-facing normal n_u at u and the outward-facing normals $n_{\tilde{x}}$ for every point $x \in \tau$ are all directed to the same side of τ. Therefore, we can assign a counterclockwise orientation to every point on τ that is consistent over τ and induces a counterclockwise ordering of τ’s vertices.

Consider the wedges W_p, $W_{p'}$, and $W_{p''}$ and their line ℓ_τ of mutual intersection, defined above. Recall that ℓ_τ is parallel to n_τ and passes through u. Recall that $\mathrm{Vor}|_{\tilde{\Sigma}}p \subset W_p$, $\mathrm{Vor}|_{\tilde{\Sigma}}p' \subset W_{p'}$, and $\mathrm{Vor}|_{\tilde{\Sigma}}p'' \subset W_{p''}$. The ordering of the sites p, p', and p'' around τ (counterclockwise or clockwise) is determined by the orientation of the normals $n_{\tilde{x}}$ relative to τ; in turn, this ordering determines the ordering of the wedges around ℓ_τ. The ordering of the cells $\mathrm{Vor}|_{\tilde{\Sigma}}p$, $\mathrm{Vor}|_{\tilde{\Sigma}}p'$, and $\mathrm{Vor}|_{\tilde{\Sigma}}p''$ around u is determined by the ordering of the wedges around ℓ_τ and the orientation of the normal n_u relative to τ. As the normals n_u and $n_{\tilde{x}}$ point to the same side of τ, the result follows. □

Consider two restricted Delaunay triangles $\tau_1 = \triangle pp'w_1$ and $\tau_2 = \triangle p'pw_2$ that satisfy the requirements of Theorem 57; hence $\nu|_{\tau_1}$ is a homeomorphism and so is $\nu|_{\tau_2}$. The two triangles share an edge pp'. The next lemma shows that their images under ν do not overlap each other; in particular, their images fall on opposite “sides” of the image of pp'.

Lemma 60. *Let $e \in \mathrm{Vor}|_{\tilde{\Sigma}}V$ be a restricted Voronoi edge with vertices u_1 and u_2, whose dual restricted Delaunay triangles are $\tau_1 = \triangle pp'w_1$ and $\tau_2 = \triangle p'pw_2$, with $p, p', w_1, w_2 \in V$. Suppose that the restricted Voronoi cells of p, p', and w_1 adjoin u_1 in counterclockwise order around u_1, and the cells of p', p, and w_2 adjoin u_2 in counterclockwise order around u_2. Let $R_1 = |pu_1| = |p'u_1| = |w_1u_1|$ and $R_2 = |pu_2| = |p'u_2| = |w_2u_2|$. Suppose that at least one of the following holds: either $R_1 \leq 0.3202\,\mathrm{lfs}(u_1)$ or $R_1 \leq 0.3606\,\mathrm{lfs}(q)$ for each $q \in \{p, p', w_1\}$. Moreover, suppose that at least one of the following holds: either $R_2 \leq 0.3202\,\mathrm{lfs}(u_2)$ or $R_2 \leq 0.3606\,\mathrm{lfs}(q)$ for each $q \in \{p, p', w_2\}$. Let $Q = \tau_1 \cup \tau_2$. Then $\nu|_Q$ is a homeomorphism from Q to its image $\nu(Q)$ on Σ.*

Proof. By Theorem 57, $\nu|_{\tau_1}$ is a continuous bijection with a continuous inverse that preserves the orientation of every projected point, and so is $\nu|_{\tau_2}$. Hence it remains only to show that $\nu|_Q$ is injective.

Suppose for the sake of contradiction that there are two distinct points $x \in \tau_1 \setminus pp'$ and $y \in \tau_2 \setminus pp'$ such that $\nu(x) = \nu(y)$. Let $\tilde{x} = \nu(x) = \nu(y)$. Let $T_{\tilde{x}}\Sigma$ be the plane tangent to Σ at $\tilde{x}$. Let $\ell_{\tilde{x}}$ be the normal segment of $\tilde{x}$, which passes through $\tilde{x}$, x, and y and is perpendicular to $T_{\tilde{x}}\Sigma$. For any point $p \in \mathbb{R}^3$, let $\bar{p}$ denote the orthogonal projection of p onto $T_{\tilde{x}}\Sigma$. (The projection direction is parallel to $\ell_{\tilde{x}}$.) By Theorem 57, neither τ_1 nor τ_2 is parallel to $\ell_{\tilde{x}}$, so the orthogonal projections of τ_1 and τ_2 onto $T_{\tilde{x}}\Sigma$ are triangles (rather than line segments). As $\tilde{x}$ lies in both orthogonal projections but not on the orthogonal projection of pp', it follows that if p, p', and w_1 occur in counterclockwise order on $T_{\tilde{x}}\Sigma$, then p', p, and w_2 occur in clockwise order on $T_{\tilde{x}}\Sigma$; and if the former sites occur in clockwise order, then the latter sites occur in counterclockwise order. Therefore, the orientation of p, p', and w_1 is opposite to the orientation of p', p, and w_2.

By Lemma 59, p, p', and w_1 occur in counterclockwise order around the boundary of τ_1; likewise, p', p, and w_2 occur in counterclockwise order around the boundary of τ_2. But this contradicts the conclusion of the previous paragraph.

Hence $\nu|_Q$ is an injection; hence $\nu|_Q$ is a bijection from Q to $\nu(Q)$. As $\nu|_{\tau_1}$ is continuous and has a continuous inverse over $\nu(\tau_1)$, and $\nu|_{\tau_2}$ is continuous and has a continuous inverse over $\nu(\tau_2)$, $\nu|_Q$ is continuous and has a continuous inverse over $\nu(Q)$. Therefore $\nu|_Q$ is a homeomorphism from Q to $\nu(Q)$. □

Lemma 61. *Consider* $\mathrm{Vor}\,|_{\tilde{\Sigma}}V$*, where* $\Sigma \subset \mathbb{R}^3$ *is a smooth* 2*-manifold without boundary and* $V \subset \Sigma$ *is a nonempty, finite sample. Suppose that for every site* $p \in V$ *and every point* $x \in \mathrm{Vor}\,|_{\tilde{\Sigma}_S}\,p$*,* $|px| < \xi\,\mathrm{lfs}(p)$*, where* $\xi = \sqrt{(\sqrt{5}-1)/2} \doteq 0.786151$*. Moreover, suppose that for every principal vertex* $u \in \mathrm{Vor}\,|_{\tilde{\Sigma}_S}\,V$*, at least one of the following holds: u is a principal vertex and* $R \leq 0.3202\,\mathrm{lfs}(u)$ *or* $R \leq 0.3606\,\mathrm{lfs}(v)$*, where* τ *is the restricted Delaunay triangle dual to u, R is the distance from u to each vertex of* τ*, and v is the vertex of* τ *at* τ*'s largest plane angle. Suppose that every restricted Voronoi vertex in* $\mathrm{Vor}\,|_{\tilde{\Sigma}}V$ *has degree three. Let T be the set of triangles in the restricted Delaunay triangulation* $\mathrm{Del}\,|_{\tilde{\Sigma}}V$*. Then the nearest point map* $\nu : |T| \to \Sigma$ *is a surjection.*

Proof. Suppose for the sake of contradiction that some point $x \in \Sigma$ is not in $\nu(|T|)$. Let $\mathring{\Sigma}$ be the connected component of Σ that contains x. By Corollary 51, there are at least six sites on $\mathring{\Sigma}$, so there is at least one principal vertex on $\mathring{\Sigma}$; let $\tau_0 \in T$ be its dual restricted Delaunay triangle. By Theorem 57, for every triangle $\tau \in T$, $\nu|_\tau$ is a homeomorphism from τ to $\nu(\tau)$, so $\nu(\tau_0)$ is a topological disk that is closed with respect to Σ. Let y be a point in the interior of $\nu(\tau_0)$ that is not in V (not a site). As $\mathring{\Sigma}$ is a connected 2-manifold, there exists a directed path $\gamma \subset \mathring{\Sigma}$ from y to x that does not intersect any site in V (except x if x is a site).

As $\nu(\tau)$ is closed with respect to Σ for every $\tau \in T$, $\nu(|T|)$ is closed with respect to Σ. Let $z \in |T|$ be a point such that γ leaves $\nu(|T|)$ for the last time at $\nu(z)$, never to re-enter. By supposition, $z \neq x$ and z is not a site. Let $\tau_1 \in T$ be a triangle that contains z. As γ leaves $\nu(|T|)$ at $\nu(z)$, γ leaves $\nu(\tau_1)$ at $\nu(z)$ and z lies on the relative interior of an edge e' of τ_1. As every restricted Voronoi vertex has degree three, there is a restricted Voronoi edge e dual to e'. One of e's vertices is dual to τ_1; let $\tau_2 \in T$ be the restricted Delaunay triangle dual to the other vertex. Let $Q = \tau_1 \cup \tau_2$. By Lemma 60, $\nu|_Q$ is a homeomorphism from Q to $\nu(Q)$, so $\nu(z)$ is in the interior of $\nu(Q)$. Therefore, where the path γ leaves $\nu(\tau_1)$ for the last time at $\nu(z)$, γ enters the interior of $\nu(\tau_2)$. This contradicts the claim that γ leaves $\nu(|T|)$ for the last time at $\nu(z)$. Therefore, for every point $x \in \Sigma$, $x \in \nu(|T|)$. □

The Nearest Point Map Is Injective

Lemma 62. *Let V be a nonempty, finite sample of* Σ*. Let p be a site in V. Let W be the set of restricted Voronoi vertices in* $\mathrm{Vor}\,|_{\tilde{\Sigma}}p$*. Let* $\dot{T} \subseteq \mathrm{Del}\,|_{\tilde{\Sigma}}V$ *be the set of restricted Delaunay triangles that* *have vertex p—that is, the set of restricted Delaunay triangles dual to the vertices in W. Suppose that for every point* $x \in \mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$, $|px| < \xi\, \mathrm{lfs}(p)$ *where* $\xi = \sqrt{(\sqrt{5}-1)/2} \doteq 0.786151$. *Moreover, suppose that for each principal vertex* $u \in W$, *at least one of the following holds:* $s \leq 0.3202\, \mathrm{lfs}(u)$ *or* $s \leq 0.3606\, \mathrm{lfs}(q)$ *for each vertex q of* τ, *where* $\tau \in \dot{T}$ *is the restricted Delaunay triangle dual to u and R is the distance from u to each vertex of* τ. *Lastly, suppose that every vertex in W has degree three.*

Then $|\dot{T}|$ *is a topological closed disk with p in its interior. Moreover, there exists an open neighborhood* $N \subset |\dot{T}|$ *of p such that* $\nu|_N$ *is a homeomorphism from N to its image* $\nu(N)$ *on* Σ.

Proof. By Theorem 49, $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$ is homeomorphic to a closed disk. By Lemma 48, the orthogonal projection of $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$ onto the tangent space $T_p\Sigma$ is star-shaped.

Choose an arbitrary axis on $T_p\Sigma$ with origin p such that each point in $T_p\Sigma \setminus \{p\}$ can be assigned an angle counterclockwise from the axis in the range $[0°, 360°)$. The rotary direction deemed "counterclockwise" is consistent with p's orientation. For convenience, we use a rotary equivalence class of angles in which θ and $\theta + 360°$ denote the same angle; so, for instance, the range $[350°, 370°]$ denotes a 20° interval of angles, proceeding counterclockwise from 350° and stopping at 10°. We extend these assigned angles to $\mathbb{R}^3$: we assign any point in $\mathbb{R}^3$ the same angle as its orthogonal projection onto $T_p\Sigma$, unless the projected point is p (in which case its angle is undefined). The star-shaped property implies that no two principal vertices of $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$ have the same angle, and that the principal vertices of $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$ sorted by increasing angle match their counterclockwise ordering around the boundary of $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$.

Each Voronoi vertex u of $\mathrm{Vor}\,|_{\overline{\Sigma}} p$ has a dual triangle $\tau \in \dot{T}$ such that, by Theorem 57, $\nu|_\tau$ is an orientation-preserving homeomorphism from τ to $\nu(\tau)$ and p's normal segment ℓ_p is not parallel to τ. Therefore, there is an angle $\psi > 0°$ such that $\angle(n_p, n_\tau) \leq 90° - \psi$ for every $\tau \in \dot{T}$. For any $\tau \in \dot{T}$, let $\bar{\tau}$ be the orthogonal projection of τ onto $T_p\Sigma$, and observe that $\bar{\tau}$ is also a triangle with straight edges. As $\angle(n_p, n_\tau) < 90°$, if the vertices p, p', and p'' of τ occur in counterclockwise order around the boundary of τ, then their projections $\bar{p}$, $\bar{p}'$, and $\bar{p}''$ on $T_p\Sigma$ occur in counterclockwise order around the boundary of $\bar{\tau}$. (Intuitively, projecting τ onto $T_p\Sigma$ does not "invert" the triangle.) If the angle assigned to p' and $\bar{p}'$ (both angles are the same) is ϕ, then the angle of p'' and $\bar{p}''$ is $\phi + \angle\bar{p}''\bar{p}\bar{p}'$, where $\angle\bar{p}''\bar{p}\bar{p}' \in (0°, 180°)$.

These facts hold not only for $T_p\Sigma$, but also if $T_p\Sigma$ is replaced by any plane whose normal vector n satisfies $\angle(n, n_p) < \psi$, because any such n also satisfies $\angle(n, n_\tau) < 90°$ for every $\tau \in \dot{T}$.

Hence the counterclockwise ordering of the principal vertices around $\mathrm{Vor}\,|_{\overline{\Sigma}_S}\, p$ implies a counterclockwise ordering of the corresponding projected triangles around p. However, it does not imply that the triangles wind only once around p; we must eliminate the possibility that the triangles wind around p two or more times.

Observe that if a restricted Voronoi vertex u is assigned an angle θ, the range of angles spanned by its dual restricted Delaunay triangle τ cannot include $\theta + 180°$, because if it did, the sphere with center u that passes through p could not enclose τ. Let $\theta_0, \theta_1, \ldots, \theta_{j-1}$ be the sorted angles of the j restricted Voronoi vertices on the boundary of $\mathrm{Vor}\,|_{\overline{\Sigma}} p$. For each such angle θ_i, let ϕ_i and ϕ_{i+1} (where the subscript $i+1$ is interpreted modulo j) be the angles of the vertices (except p) of the corresponding dual restricted Delaunay triangle. For the triangles in $\dot{T}$ to wind around p two or more times, there must be an i such that $\theta_i + 180° \in [\phi_i, \phi_{i+1}]$. As this is impossible, the triangles wind around p only once.

Each projected triangle is confined to its own range of angles $[\phi_i, \phi_{i+1}]$, and the only points it shares with other triangles lie on the shared edges at angles ϕ_i and ϕ_{i+1} and on the shared vertex p. Therefore, the projection of $|\dot{T}|$ onto $T_p\Sigma$ is an injection, and $|\dot{T}|$ is a topological closed disk with p in its interior.

To show that ν restricted to some open neighborhood $N \subset |\dot{T}|$ of p is a homeomorphism, we take advantage of the fact that for every plane Π whose normal vector n satisfies $\angle(n, n_p) < \psi$, the orthogonal projection of $|\dot{T}|$ onto Π is an injection. Therefore, every line $\ell \subset \mathbb{R}^3$ such that $\angle(\ell, n_p) < \psi$ intersects $|\dot{T}|$ in at most one point. We choose $N = |\dot{T}| \cap B$ where B is an open ball centered at p and small enough that for every point $q \in \nu(N)$, $\angle(n_q, n_p) < \psi$. A sufficiently small ball satisfies this condition, because ν is continuous over $|\dot{T}|$ and n_q is continuous for $q \in \Sigma$, so their composition is continuous. The condition guarantees that $\nu|_N$ is injective; hence $\nu|_N$ is a bijection from N to $\nu(N)$. As $\nu|_N$ is continuous and has a continuous inverse over $\nu(N)$, $\nu|_N$ is a homeomorphism from N to $\nu(N)$. □

Lemma 63. *Consider* $\mathrm{Vor}|_{\tilde{\Sigma}}V$, *where* $\Sigma \subset \mathbb{R}^3$ *is a smooth* 2*-manifold without boundary and* $V \subset \Sigma$ *is a nonempty, finite sample. Suppose that for every site* $p \in V$ *and every point* $x \in \mathrm{Vor}|_{\tilde{\Sigma}_S} p$, $|px| < \xi\,\mathrm{lfs}(p)$, *where* $\xi = \sqrt{(\sqrt{5} - 1)/2} \doteq 0.786151$. *Moreover, suppose that for every principal vertex* $u \in \mathrm{Vor}|_{\tilde{\Sigma}}V$, *at least one of the following holds: u is a principal vertex and* $s \leq 0.3202\,\mathrm{lfs}(u)$, *or* $s \leq 0.3606\,\mathrm{lfs}(v)$, *where* τ *is the restricted Delaunay triangle dual to u, R is the distance from u to each vertex of* τ, *and v is the vertex of* τ *at* τ*'s largest plane angle. Suppose that every restricted Voronoi vertex of* $\mathrm{Vor}|_{\tilde{\Sigma}}V$ *has degree three. Then the nearest point map* $\nu : |\mathrm{Del}|_{\tilde{\Sigma}}V| \to \Sigma$ *is an injection.*

Proof. Suppose for the sake of contradiction that there are two distinct points $x, y \in |\mathrm{Del}|_{\tilde{\Sigma}}V|$ such that $\nu(x) = \nu(y)$. By Corollary 51, there are at least six sites on each connected component of Σ, so there are at least two restricted Voronoi vertices in each restricted Voronoi cell; hence every restricted Delaunay vertex and every restricted Delaunay edge is a subset of some restricted Delaunay triangle. It follows that x lies on some restricted Delaunay triangle, and likewise for y.

Let T be the set of triangles in the restricted Delaunay triangulation $\mathrm{Del}|_{\tilde{\Sigma}}V$. Let $\tau_x, \tau_y \in T$ be triangles containing x and y, respectively. Theorem 57 implies that no triangle in T contains both x and y; one implication is that $\tau_x \neq \tau_y$. By Corollary 53, no site intersects $\nu(\tau_x)$ except the three vertices of τ_x; hence y is not a site. Symmetrically, x is not a site.

Let p be a vertex of τ_x not shared by τ_y. Let $T_p \subset T$ be the set of restricted Delaunay triangles that have p for a vertex (including τ_x). Let $\gamma \subset \nu(\tau_x)$ be a directed path on Σ from $\nu(x) = \nu(y)$ to p such that $\gamma \setminus \{\nu(x), p\}$ lies in the relative interior of $\nu(\tau_x)$. By Lemma 62, there exists an open neighborhood $N \subset |T_p|$ of p such that the nearest point map $\nu|_N$ is a homeomorphism from N to its image $\nu(N)$ on Σ. By Corollary 53, p does not intersect $\nu(\tau)$ for any restricted Delaunay triangle $\tau \in T \setminus T_p$, so we can assume without loss of generality that N is sufficiently small that $\nu(N)$ does not intersect $\nu(\tau)$ for any triangle $\tau \in T \setminus T_p$. Let $w \in \tau_x$ be a point such that $\nu(w) \in \gamma \cap \nu(N) \setminus \{\nu(x), p\}$.

Observe that w is in the relative interior of τ_x, because $\gamma \setminus \{\nu(x), p\}$ is a subset of the relative interior of $\nu(\tau_x)$.

Let $T_{\neg x} = T \setminus \{\tau_x\}$, the set containing all the restricted Delaunay triangles except τ_x. We claim that there is a point $w' \in |T_{\neg x}|$ such that $\nu(w') = \nu(w)$. We establish the claim with essentially the same method used to prove Lemma 61. Suppose for the sake of contradicting this claim that $\nu(w) \notin \nu(|T_{\neg x}|)$. Observe that the path γ starts at $\nu(y)$, which is a subset of $\nu(|T_{\neg x}|)$ because $\tau_y \in T_{\neg x}$. Hence there exists a point $z \in |T_{\neg x}|$ such that γ leaves $\nu(|T_{\neg x}|)$ at $\nu(z)$ for the last time before reaching $\nu(w)$. Let $\tau_1 \in T_{\neg x}$ be a triangle that contains z. As γ leaves $\nu(|T_{\neg x}|)$ at $\nu(z)$, γ leaves $\nu(\tau_1)$ at $\nu(z)$, so z lies on the relative interior of an edge e' of τ_1. Let e be the restricted Voronoi edge dual to e'. One of e's vertices is dual to τ_1; let $\tau_2 \in T$ be the restricted Delaunay triangle dual to the other vertex. As $\nu(z)$ lies on γ in the relative interior of $\nu(\tau_x)$ and $\nu(z)$ also lies on $\nu(e')$, e' is not an edge of τ_x. Therefore, $\tau_2 \neq \tau_x$ and $\tau_2 \in T_{\neg x}$. By Lemma 60, there is an open neighborhood $N_z \subset \tau_1 \cup \tau_2$ of z such that $\nu(N_z)$ is a neighborhood of $\nu(z)$ that is open with respect to Σ. Therefore, where the path γ leaves $\nu(\tau_1)$ for the last time at $\nu(z)$, γ remains in $\nu(N_z)$ a little further, so it enters the interior of $\nu(\tau_2)$. This contradicts the fact that γ leaves $\nu(|T_{\neg x}|)$ for the last time at $\nu(z)$. Hence, $\nu(w) \in \nu(|T_{\neg x}|)$; that is, there is a triangle $\tau_{w'} \in T_{\neg x}$ and a point $w' \in \tau_{w'}$ such that $\nu(w') = \nu(w)$.

As $\nu(w') \in \nu(N) \cap \nu(\tau_{w'})$ and $\nu(N)$ does not intersect $\nu(\tau)$ for any restricted Delaunay triangle $\tau \notin T_p$, $\tau_{w'} \in T_p$. But $\tau_{w'} \neq \tau_x$. The triangles in T_p intersect only at p and along their shared edges (because they are triangles in a three-dimensional Delaunay triangulation), so $\tau_{w'}$ does not intersect the relative interior of τ_x. As $w' \in \tau_{w'}$ and w is in the relative interior of τ_x, $w' \neq w$. Hence there exist two distinct points $w, w' \in N$ such that $\nu(w) = \nu(w')$, contradicting the fact that $\nu|_N$ is a homeomorphism. By this contradiction, we conclude that there are no two distinct points $x, y \in |\,\mathrm{Del}\,|_{\tilde{\Sigma}} V|$ such that $\nu(x) = \nu(y)$, and therefore $\nu : |\,\mathrm{Del}\,|_{\tilde{\Sigma}} V| \to \Sigma$ is an injection. □

Theorem 64. *Consider* $\mathrm{Vor}\,|_{\tilde{\Sigma}} V$, *where* $\Sigma \subset \mathbb{R}^3$ *is a smooth* 2*-manifold without boundary and* $V \subset \Sigma$ *is a nonempty, finite sample. Suppose that for every site* $p \in V$ *and every point* $x \in \mathrm{Vor}\,|_{\tilde{\Sigma}_S}\, p$, $|px| < \xi\, \mathrm{lfs}(p)$, *where* $\xi = \sqrt{(\sqrt{5} - 1)/2} \doteq 0.786151$. *Suppose that every restricted Voronoi vertex in* $\mathrm{Vor}\,|_{\tilde{\Sigma}} V$ *has degree three. Moreover, suppose that for every principal vertex* $u \in \mathrm{Vor}\,|_{\tilde{\Sigma}} V$, $R \leq 0.3202\,\mathrm{lfs}(u)$, *and for every secondary vertex* $u \in \mathrm{Vor}\,|_{\tilde{\Sigma}} V$, *u's dual restricted Delaunay triangle* τ *satisfies* $R \leq 0.3606\,\mathrm{lfs}(q)$ *for each vertex* q *of* τ, *where* R *is the distance from* u *to each vertex of* τ. *Then the nearest point map* $\nu : |\,\mathrm{Del}\,|_{\tilde{\Sigma}} V| \to \Sigma$ *is a homeomorphism.*

Proof. By Lemma 61, $\nu : |\,\mathrm{Del}\,|_{\tilde{\Sigma}} V| \to \Sigma$ is a surjection. By Lemma 63, it is an injection too. Hence it has an inverse defined over Σ. Both the nearest point map and its inverse are continuous, so it is a homeomorphism. □

This brings us to our main result, Theorem 23. Recall its statement:

Let V be a constrained ϵ-sample of (Σ, S, Z) for some $\epsilon \leq 0.3202$. Suppose that for every segment $pq \in S$, $|pq| \leq 0.3368\,\mathrm{lfs}(p)$. Suppose that every extended Voronoi vertex in $\mathrm{Vor}\,|_{\tilde{\Sigma}} V$ has degree three. Suppose that for every extended Voronoi vertex that lies on an extrusion, its dual restricted Delaunay triangle satisfies $r \leq 0.3606\,\mathrm{lfs}(v)$, where r is τ's circumradius and v is the vertex of τ at τ's largest plane angle. Then the nearest point map $\nu : |\,\mathrm{Del}\,|_{\tilde{\Sigma}} V| \to \Sigma$ is a homeomorphism.

Proof. For every point $x \in \overline{\Sigma}_S$, the nearest site $p \in V$ satisfies $|px| \leq 0.3202\,\mathrm{lfs}(x)$ by the definition of constrained ϵ-sample. By the Feature Translation Lemma (Lemma 38), $|px| \leq 0.3202/(1 - 0.3202)\,\mathrm{lfs}(p) < 0.48\,\mathrm{lfs}(p)$. Therefore, for every site $p \in V$ and every point $x \in \mathrm{Vor}|_{\overline{\Sigma}_S}\, p$, $|px| < \xi\,\mathrm{lfs}(p)$, satisfying one of the conditions of Theorem 64.

Let $u \in \mathrm{Vor}|_{\overline{\Sigma}} V$ be a restricted Voronoi vertex and let $\tau \in \mathrm{Del}|_{\overline{\Sigma}} V$ be the restricted Delaunay triangle dual to u. If u lies on the principal surface $\overline{\Sigma}_S$, let s be the distance from u to any vertex of τ. As τ's vertices are the sites closest to u (that are visible from u), for every vertex q of τ, $s = |qu| < 0.3202\,\mathrm{lfs}(u)$, which satisfies a condition of Theorem 64. If u lies on an extrusion, let r be τ's circumradius and let v be the vertex of τ at τ's largest plane angle. By assumption, $r \leq 0.3606\,\mathrm{lfs}(v)$, which satisfies the remaining condition of Theorem 64.

By Theorem 64, $\nu : |\mathrm{Del}|_{\overline{\Sigma}} V| \to \Sigma$ is a homeomorphism. □

Appendix B

Proofs for Chapter 5

B.1 Proof of Theorem 35

The proof of Theorem 35 is the combination of the following lemmas.

Lemma 65. *Let $w \in \mathbb{R}^d$ be any vector and let $w_{||}$ be the orthogonal projection of w onto* rowspace(X). *Then, for the objective function*

$$\mathcal{L}_2(X, y; w) = \frac{1}{2}||Xw - y||_2^2 + \epsilon||w||_2||Xw - y||_1 + \frac{\epsilon^2 n}{2}||w||_2^2.$$

we have that $\mathcal{L}_2(X, y; w) \geq \mathcal{L}_2(X, y; w_{||})$, with equality if and only if $w = w_{||}$. Hence for any optimal solution w^* of $\mathcal{L}_2$, $w^* \in$ rowspace(X).

Proof. Let $w = w_{||} + w_{\perp}$ be *any* vector $\mathbb{R}^d$ where $w_{||} \in$ rowspace(X) and $w_{\perp} \in$ nullspace(X).

$$\begin{aligned}
\mathcal{L}_2(X, y; w) &= \frac{1}{2}||Xw - y||_2^2 + \epsilon||w||_2||Xw - y||_1 + \frac{\epsilon^2 n}{2}||w||_2^2 \\
&= \frac{1}{2}||X(w_{||} + w_{\perp}) - y||_2^2 + \epsilon||w_{||} + w_{\perp}||_2||X(w_{||} + w_{\perp}) - y||_1 + \frac{\epsilon^2 n}{2}||w_{||} + w_{\perp}||_2^2 \\
&= \frac{1}{2}||Xw_{||} - y||_2^2 + \epsilon||w_{||} + w_{\perp}||_2||Xw_{||} - y||_1 + \frac{\epsilon^2 n}{2}||w_{||} + w_{\perp}||_2^2 \\
&= \frac{1}{2}||Xw_{||} - y||_2^2 + \epsilon\sqrt{||w_{||}||_2^2 + ||w_{\perp}||_2^2}||Xw_{||} - y||_1 + \frac{\epsilon^2 n}{2}\left(||w_{||}||_2^2 + ||w_{\perp}||_2^2\right) \\
&\geq \frac{1}{2}||Xw_{||} - y||_2^2 + \epsilon||w_{||}||_2||Xw_{||} - y||_1 + \frac{\epsilon^2 n}{2}||w_{||}||_2^2
\end{aligned}$$

with equality if and only if $||w_{\perp}|| = 0$. The third equality follows from the fact that $w_{\perp}$ is in nullspace(X), the fourth from the fact that $w_{||} \perp w_{\perp}$. This proves the first statement. The second statement regarding w^* follows immediately. □

Lemma 66. *Let $C \in \mathcal{H}$ be a convex cell with signature s. The restriction of $\mathcal{L}_2$ to the interior of C, denoted $\mathcal{L}_2|_{\operatorname{Int} C}$, is a convex function. Furthermore, if $s \neq -y$ then $\mathcal{L}_2|_{\operatorname{Int} C}$ is a strongly convex function.*

Suppose that $s = -y$, meaning that C contains the origin. There are four possible cases, three of which depend on the value of ϵ.

1. *If $Xw = y$ is an inconsistent system, then $\mathcal{L}_2|_{\operatorname{Int} C}$ is a strongly convex function.*

2. *If $Xw = y$ is a consistent system and $\epsilon \in (0, \frac{1}{\|X^\dagger y\|_2})$ then $\mathcal{L}_2|_{\operatorname{Int} C}$ is a convex function. Specifically, $\mathcal{L}|_{\operatorname{Int} C}$ is convex but not strongly convex along two line segments both of which have one endpoint at the origin and terminate at $X^\dagger y \pm u$ for some $u \in \operatorname{nullspace}(X)$ respectively. The gradient at every point on these line segments is nonzero, and so the optimal solution is found in the* rowspace(X) *at a point of strong convexity.*

3. *If $Xw = y$ is a consistent system and $\epsilon = \frac{1}{\|X^\dagger y\|_2}$, then $\mathcal{L}_2|_{\operatorname{Int} C}$ is a convex function. Specifically $\mathcal{L}|_{\operatorname{Int} C}$ is convex but not strongly convex along a single line segment with one endpoint at the origin and the other endpoint at $X^\dagger y$. The optimal solution may lie along this line.*

4. *If $Xw = y$ is a consistent system and $\epsilon > \frac{1}{\|X^\dagger y\|_2}$, then $\mathcal{L}_2|_{\operatorname{Int} C}$ is a strongly convex function.*

Proof. Let C be any cell in the hyperplane arrangement induced by $\|Xw - y\|_1$ and let $s \in \pm 1^n$ denote the signature of C. We will show that the Hessian matrix within C is positive semi-definite.

The Hessian matrix $H(w)$ at a point $w \in \operatorname{Int} C$ is

$$X^\top X + \frac{\epsilon}{\|w\|_2}(X^\top s w^\top + w s^\top X) + \frac{\epsilon^2 n}{\|w\|_2^2} w w^\top - \frac{\epsilon}{\|w\|_2^3} s^\top (Xw - y + \epsilon \|w\|_2 s) w w^\top + \frac{\epsilon}{\|w\|_2} s^\top (Xw - y + \epsilon \|w\|_2 s) I.$$

This form of the Hessian comes from twice differentiating Equation 5.15 and is equivalent to twice differentiating Equation 5.16. Note that it is crucial in the third term that the sign function is always ±1 and not defined as 0 when the input is 0; this is from where the factor of n is derived. At a high level, we examine the curvature induced by $H(w)$ in each unit direction $v \in \mathbb{S}^{d-1}$ at w and show that it is everywhere non-negative. It is worth taking a moment to examine how each term of $H(w)$ affects the curvature of the objective at w.

The term $X^\top X$ is a positive semi-definite matrix and induces a quadratic form with positive curvature in each eigen-direction whose corresponding eigenvalue is positive, and zero curvature in every eigen-direction corresponding to a zero eigenvalue.

The term $\frac{1}{\|w\|_2}(X^\top s w^\top + w s^\top X)$ is a sum of outer-product matrices. Note that this matrix is symmetric, since $(X^\top s w^\top)^\top = w s^\top X$. This matrix has a $(d-2)$-dimensional nullspace, corresponding to the intersection nullspace$(w) \cap$ nullspace$(X^\top s)$. On the 2-dimensional subspace spanned by $\{\frac{w}{\|w\|_2}, X^\top s\}$, and with respect to that basis, the outer-product has the matrix

$$\begin{pmatrix} \frac{w}{\|w\|_2} \cdot X^\top s & \frac{w}{\|w\|_2} \cdot \frac{w}{\|w\|_2} \\ X^\top s \cdot X^\top s & \frac{w}{\|w\|_2} \cdot X^\top s \end{pmatrix}.$$

The eigenvalues, within this subspace, are

$$\frac{w}{\|w\|_2} \cdot (X^\top s) \pm \sqrt{\left(\frac{w}{\|w\|_2} \cdot \frac{w}{\|w\|_2}\right)((X^\top s) \cdot (X^\top s))} = \frac{w}{\|w\|_2} \cdot (X^\top s) \pm \|X^\top s\|_2.$$

By triangle inequality, one of these eigenvalues is always positive while the other is always negative. Thus there is one direction of positive curvature and one direction of negative curvature. The eigenvectors are

$$\frac{1}{\sqrt{1 + \|X^\top s\|_2^2}} X^\top s \pm \frac{\|X^\top s\|_2}{\sqrt{1 + \|X^\top s\|_2^2}} \frac{w}{\|w\|_2}.$$

The term $\frac{\epsilon^2 n}{\|w\|_2^2} ww^\top$ induces positive curvature in the direction w with eigenvalue $\epsilon^2 n$ and 0 curvature in every direction orthogonal to w.

The term $-\frac{\epsilon}{\|w\|_2^3} s^\top(Xw - y + \epsilon\|w\|_2 s)ww^\top$ induces negative curvature in the direction w with eigenvalue $-\frac{\epsilon}{\|w\|_2} s^\top(Xw - y + \epsilon\|w\|_2 s)$. However the negative curvature in the direction w is exactly undone by the positive curvature induced by the term $\frac{\epsilon}{\|w\|_2} s^\top(Xw - y + \epsilon\|w\|_2 s)I$ which induces positive curvature in every direction with eigenvalues all equal to $\frac{\epsilon}{\|w\|_2} s^\top(Xw - y + \epsilon\|w\|_2 s)$. The result of the sum of these two terms is a quadratic form which induces 0 curvature in the direction w and positive curvature in every direction orthogonal to w with eigenvalue $\frac{\epsilon}{\|w\|_2} s^\top(Xw - y + \epsilon\|w\|_2 s)$. Note that the value $\frac{\epsilon}{\|w\|_2} s^\top(Xw - y + \epsilon\|w\|_2 s)$ is positive by definition, since w is in the convex cell with signature s, and so

$$\frac{\epsilon}{\|w\|_2} s^\top(Xw - y + \epsilon\|w\|_2 s) = \frac{\epsilon}{\|w\|_2}(\|Xw - y\|_1 + \epsilon\|w\|_2 n) > 0.$$

Let $v \in \mathbb{S}^{d-1}$ be a unit vector. The curvature in the direction v is proportional (with positive constant of proportionality) to

$$\begin{aligned} v^\top H(w) v &= v^\top X^\top X v + \frac{\epsilon}{\|w\|_2}\left(v^\top (X^\top s w^\top + w s^\top X) v\right) + \frac{\epsilon^2 n}{\|w\|_2^2} v^\top w w^\top v \\ &\quad - \frac{\epsilon}{\|w\|_2^3} s^\top (Xw - y + \epsilon\|w\|_2 s) v^\top w w^\top v + \frac{\epsilon}{\|w\|_2} s^\top(Xw - y + \epsilon\|w\|_2 s) v^\top v \\ &= \|Xv\|_2^2 + \frac{2\epsilon}{\|w\|_2}(w^\top v)(s^\top X v) + \frac{\epsilon^2 n}{\|w\|_2^2}(w^\top v)^2 + \epsilon\left(\frac{\|Xw - y\|_1}{\|w\|_2} + \epsilon n\right)\left(1 - \left(\frac{w}{\|w\|_2} \cdot v\right)^2\right) \\ &= \underbrace{\|Xv\|_2^2 + \frac{2\epsilon\sqrt{n}}{\|w\|_2}(w^\top v)\|Xv\|_2 \cos\varphi + \frac{\epsilon^2 n}{\|w\|_2^2}(w^\top v)^2}_{\text{term 1}} + \underbrace{\epsilon\left(\frac{\|Xw - y\|_1}{\|w\|_2} + \epsilon n\right)\left(1 - \cos^2\theta\right)}_{\text{term 2}} \end{aligned}$$

where $\varphi = \angle(s, Xv)$ and $\theta = \angle(w, v)$. It's easy to see that term 2 is always greater than or equal to 0, since $\cos^2\theta \in [0, 1]$, with equality when $\cos^2\theta = 1$. By the quadratic formula, term 1 is also always greater than or equal to 0, with equality when $\cos\varphi = \pm 1$ *and* $\text{sign}(\cos\varphi) \neq \text{sign}(w^\top v)$; otherwise the zeros given by the quadratic formula have an imaginary component that depends on $\sin\varphi$. Thus, at this point, we see that $H(w)$ is at least positive semi-definite in $\operatorname{Int} C$.

We wish to derive under which conditions this inequality is strict, implying that $H(w)$ is positive-definite in C. First we will show that if w is in a cell of the hyperplane arrangement whose signature is $s \neq -y$, then $H(w)$ is positive definite. The conditions which must be true for $v^\top H(w)v = 0$ imply that $s = -y$; w must be in the cell that contains the origin.

For $v^\top H(w)v = 0$ we need *both* term 1 *and* term 2 to be equal to 0. Term 2 is equal to 0 if and only if $\cos\theta = \pm 1$, which implies that $v \parallel w$. Since v is a unit vector, we have $v = \pm\frac{w}{\|w\|_2}$. For term 1 to be equal to 0 we need $\cos\varphi = \pm 1$ and $\operatorname{sign}(\cos\varphi) \neq \operatorname{sign}(w^\top v)$. The first of these two conditions implies that $s \parallel Xv$. Suppose that $v = \frac{w}{\|w\|_2}$; then $Xv = -\alpha s$ for some $\alpha > 0$. So we have that

$$\begin{aligned} \frac{Xv}{\|Xv\|_2} &= -\frac{s}{\|s\|_2} \\ \frac{Xw}{\|Xw\|_2} &= \\ Xw &= -\frac{\|Xw\|_2}{\|s\|_2} s \\ x_i^\top w &= -\frac{\|Xw\|_2}{\|s\|_2} s_i \end{aligned}$$

Now, $w \in C$, which implies that $\operatorname{sign}(x_i^\top w - y_i) = s_i$. If $s_i = 1$, then

$$\begin{aligned} x_i^\top w - y_i &> 0 \\ x_i^\top w &> y_i \\ -\frac{\|Xw\|_2}{\|s\|_2} s_i &> y_i \\ 0 > -\frac{\|Xw\|_2}{\|s\|_2} s_i &> y_i \end{aligned}$$

which implies that $y_i = -1$. The case where $s_i = -1$ is similar, as is the case where $v = -\frac{w}{\|w\|_2}$. All together, we have that $s = -y$.

Thus the necessary (*not* sufficient) conditions for $v^\top H(w)v = 0$ can only be satisfied if $s = -y$. If $s \neq -y$ then $H(w)$ is defined and positive-definite everywhere in $\operatorname{Int} C$.

We now turn our attention toward a necessary condition, which when combined with our other necessary conditions, give a set of sufficient conditions for $H(w)$ to be positive semi-definite but not positive-definite. Suppose that $\cos\theta = \pm 1, \cos\varphi = \pm 1$ and $\operatorname{sign}(\cos\varphi) \neq \operatorname{sign}(w^\top v)$. By the above discussion $s = -y$. Suppose $v = \frac{w}{\|w\|_2}$ and, thus, $\cos\varphi = -1$. Under these conditions we have that

$$
\begin{aligned}
v^\top H(w)v &= \|Xv\|_2^2 - \frac{2\epsilon\sqrt{n}}{\|w\|_2}(w^\top v)\|Xv\|_2 + \frac{\epsilon^2 n}{\|w\|_2^2}(w^\top v)^2 + \epsilon\left(\frac{\|Xw-y\|_1}{\|w\|_2} + \epsilon n\right)\left(1-\cos^2\theta\right) \\
&= \left(\|Xv\|_2 - \frac{\epsilon\sqrt{n}}{\|w\|_2}w^\top v\right)^2 \\
&= \left(\|Xv\|_2 - \epsilon\sqrt{n}\right)^2.
\end{aligned}
$$

Note that when any one of the conditions detailed in the previous paragraph do not hold, the first equality is instead a lower bound on $v^\top H(w)v$. From this we see that the final necessary condition for $v^\top H(w)v = 0$ is for $\|Xv\|_2 = \epsilon\sqrt{n}$. Since $\cos\varphi = -1$, we must have $Xv = -\epsilon s = \epsilon y$ which implies $v = \epsilon X^\dagger y + u$ for $u \in \text{nullspace}(X)$. Recall that v is a unit vector, so $\|v\|_2^2 = \|\epsilon X^\dagger y + u\|_2^2 = \epsilon^2\|X^\dagger y\|_2^2 + \|u\|_2^2 = 1$, from which it follows that $\epsilon = \frac{\sqrt{1-\|u\|_2^2}}{\|X^\dagger y\|_2}$.

The relationship $\epsilon = \frac{\sqrt{1-\|u\|_2^2}}{\|X^\dagger y\|_2}$ gives three intervals for ϵ in which the curvature of $\mathcal{L}_2$ behaves qualitatively differently. For $\epsilon \in (0, 1/\|X^\dagger y\|_2)$ the equation has two solutions $\pm u$ in the nullspace(X) with $\|u\|_2 < 1$. Since $w \parallel v$ this ray of 0 curvature lies outside of rowspace(X), and, by Lemma 65, the gradient cannot be 0 along this ray. For $\epsilon = 1/\|X^\dagger y\|_2$, the solution is given by $u = 0$ and so there is a single ray in the direction of $X^\dagger y$ in the rowspace(X). This ray is parameterized by $\alpha X^\dagger y$ for $\alpha \in (0, 1)$. The gradient may or may not be zero along this ray. Finally for $\epsilon > 1/\|X^\dagger y\|_2$ there is no solution to the relationship and $\mathcal{L}_2$ is strongly convex within C with signature $s = -y$.

Before concluding we must address the fact that the Hessian H is not defined at $w = 0$. Let $\{O_i\}_i$ be the set of 2^d closed orthants of $\mathbb{R}^d$. We further subdivide C with signature $s = -y$ as $C_i = C \cap O_i$. Within the relative interiors of each C_i, $\mathcal{L}_2|_{\text{Int}\,C_i}$ is twice differentiable everywhere with Hessian as described above. Thus $\mathcal{L}_2|_{\text{Int}\,C_i}$ is convex for all i.

Let $w \in \text{Bd}\,C_i \cap \text{Int}\,C$ and $w \neq 0$. Then the subdifferential $\partial\mathcal{L}_2|_{C_i}(w)$ is nonempty and, in particular, contains the gradient $\nabla\mathcal{L}_2(w)$, which is defined at w since $w \in \text{Int}\,C$ and $w \neq 0$. The intersection $\partial\mathcal{L}_2|_{C_i}(w) \cap \partial\mathcal{L}_2|_{C_j}(w) = \{\nabla\mathcal{L}_2(w)\}$ for $w \in \text{Bd}\,C_i \cap \text{Bd}\,C_j \cap \text{Int}\,C$, since $\mathcal{L}_2$ is actually differentiable at w.

Now let $w \in \text{Int}\,C_i$ and $w' \in \text{Int}\,C_j$ for $i \neq j$ and such that the line segment ww' does not intersect the origin in its relative interior. Further suppose that C_i and C_j are adjacent along the line segment ww', meaning that there is a single point w'' at which the line segment ww' leaves $\text{Int}\,C_i$ and enters

$\operatorname{Int} C_j$. Note that w, w'', w' are collinear. Then

$$\begin{aligned}
\langle \nabla \mathcal{L}_2(w), w' - w\rangle &= \langle \nabla \mathcal{L}_2(w), w' - w''\rangle + \langle \nabla \mathcal{L}_2(w), w'' - w\rangle \\
&= \frac{\|w'' - w\|_2}{\|w' - w''\|_2}\langle \nabla \mathcal{L}_2(w), w'' - w\rangle + \langle \nabla \mathcal{L}_2(w), w'' - w\rangle \\
&\leq \frac{\|w'' - w\|_2}{\|w' - w''\|_2}\langle \nabla \mathcal{L}_2(w''), w'' - w\rangle + \langle \nabla \mathcal{L}_2(w), w'' - w\rangle \\
&= \langle \nabla \mathcal{L}_2(w''), w' - w''\rangle + \langle \nabla \mathcal{L}_2(w), w'' - w\rangle \\
&\leq \mathcal{L}_2|_{\operatorname{Int} C_j}(w') - \mathcal{L}_2|_{\operatorname{Int} C_j}(w'') + \mathcal{L}_2|_{\operatorname{Int} C_i}(w'') - \mathcal{L}_2|_{\operatorname{Int} C_i}(w) \\
&= \mathcal{L}_2|_{\operatorname{Int} C_j}(w') - \mathcal{L}_2|_{\operatorname{Int} C_i}(w) \\
&= \mathcal{L}_2(w') - \mathcal{L}_2(w).
\end{aligned}$$

The second equality follows from collinearity and the first inequality follows from convexity of $\mathcal{L}_2|_{\operatorname{Int} C_i}$ from which we can derive $\langle \nabla \mathcal{L}_2(w'') - \nabla \mathcal{L}_2(w), w'' - w\rangle \geq 0$. The remaining steps are straightforward. This argument can be extended to a line segment ww' for w and w' in two cells that only intersect at the origin in a straightforward manner using induction. Thus it follows that $\mathcal{L}_2|_{\operatorname{Int} C}$ is convex along ww'. All that remains is the case where ww' intersects the origin.

Suppose that ww', parameterized by $\ell(t) = (1-t)w + tw'$ for $t \in [0, 1]$, intersects the origin. Choose any unit vector v such that v is not parallel to ww'. Then consider the perturbed line segment $\tilde{\ell}(t, \epsilon) = (1-t)(w + \epsilon v) + t(w' + \epsilon v) = \ell(t) + \epsilon v$ for $\epsilon > 0$. Let t_0 be such that $\ell(t_0) = 0$. As $\epsilon \to 0$, $\tilde{\ell}(t, \epsilon) \to \ell(t)$ and, in particular, $\tilde{\ell}(t_0, \epsilon) \to 0$. Since v is not parallel with ww', $\tilde{\ell}(t, \epsilon)$ does not intersect the origin for $\epsilon > 0$, and so $\mathcal{L}_2|_{\operatorname{Int} C}(\tilde{\ell}(t, \epsilon)) \leq (1-t)\mathcal{L}_2|_{\operatorname{Int} C}(\tilde{\ell}(0, \epsilon)) + t\mathcal{L}_2|_{\operatorname{Int} C}(\tilde{\ell}(1, \epsilon))$. Taking $\epsilon \to 0$ convexity follows from continuity of $\mathcal{L}_2|_{\operatorname{Int} C}$. Note that this approach only applies when $d \geq 2$; however the $d = 1$ for $\mathcal{L}_2$ case is straightforward.

□

Lemma 67. *$\mathcal{L}_2$ is a convex function. If $\mathcal{L}_2|_{\operatorname{Int} C}$ for C with signature $s = -y$ is a strongly convex function, then $\mathcal{L}_2$ is a strictly convex function. Furthermore transitions between two cells are strictly convex.*

Proof. Let $w, w' \in \mathbb{R}^d$ be any two points. The line segment ww' with endpoints w and w' is parameterized by $w_t = (1-t)w + tw'$ for $t \in [0, 1]$. If $ww' \subset \operatorname{Int} C$ for some C then Lemma 66 gives the results. Suppose that $w \in C$ and $w' \in C'$ are in distinct cells of the hyperplane arrangement and that ww' intersect the boundaries of these cells at $t_1, \ldots, t_m$. This partitions the interval $[0, 1]$ into $m + 1$ subintervals $[0, t_1] \cup [t_1, t_2] \cup \ldots \cup [t_m, 1]$, in each of which the function $\mathcal{L}_2$ is convex along $w_{t_i} w_{t_{i+1}}$ by Lemma 66.

Consider the base case where $m = 1$. The point $w_1 \in C \cap C'$, where the line segment ww' leaves C and enters C'. The facet $f = C \cap C'$ is a $(d-k)$-dimensional facet, where k is the number of hyperplanes that intersect at w_1. Said differently, at w_1 the signs of k hyperplane equations $x_i^\top w - y_i$ flip.

Imagine removing these k hyperplanes, then w and w' lie in the same cell of the induced hyperplane arrangement, and, by Lemma 66, the objective function $\mathcal{L}_2^{(-k)}$ with these k hyperplanes removed is convex. (Simply repeat the argument for $n - k$ samples.) Thus we have

$$
\begin{aligned}
\mathcal{L}_2(w_{t_1}) &= \mathcal{L}_2^{(-k)}(w_{t_1}) + \frac{1}{2}\sum_i (\langle x_i, w_{t_1}\rangle - y_i + \epsilon \operatorname{sign}(\langle x_i, w_{t_1}\rangle - y_i)\|w_{t_1}\|_2)^2 \\
&= \mathcal{L}_2^{(-k)}(w_{t_1}) + \frac{1}{2}\sum_i (\epsilon \operatorname{sign}(\langle x_i, w_{t_1}\rangle - y_i)\|w_{t_1}\|_2)^2 \\
&= \mathcal{L}_2^{(-k)}(w_{t_1}) + \frac{1}{2}\sum_i \epsilon^2 \|w_{t_1}\|_2^2 \\
&\le (1-t_1)\mathcal{L}_2^{(-k)}(w) + t_1 \mathcal{L}_2^{(-k)}(w') + (1-t_1)\frac{1}{2}\sum_i \epsilon^2\|w\|_2^2 + t_1 \frac{1}{2}\sum_i \epsilon^2 \|w'\|_2^2 \\
&< (1-t_1)\mathcal{L}_2^{(-k)}(w) + t_1 \mathcal{L}_2^{(-k)}(w') \\
&\qquad + (1-t_1)\frac{1}{2}\sum_i \Big((\langle x_i, w\rangle - y_i)^2 + 2\epsilon\|w\|_2 \operatorname{sign}(\langle x_i, w\rangle - y_i)\,(\langle x_i, w\rangle - y_i) + \epsilon^2\|w\|_2^2\Big) \\
&\qquad + t_1\frac{1}{2}\sum_i \Big((\langle x_i, w'\rangle - y_i)^2 + 2\epsilon\|w'\|_2 \operatorname{sign}(\langle x_i, w'\rangle - y_i)\,(\langle x_i, w'\rangle - y_i) + \epsilon^2\|w'\|_2^2\Big) \\
&= (1-t_1)\mathcal{L}(w) + t_1\mathcal{L}(w')
\end{aligned}
$$

The second equality follows from the crucial fact that, at w_{t_1}, each hyperplane constraint $x_i^\top w - y_i = 0$. The first inequality follows from the convexity of $\mathcal{L}_2^{(-k)}$ and $\|w\|_2$. The second inequality follows from adding strictly positive terms. The final equality follows by definition. With this fact we are ready to show the convexity of $\mathcal{L}_2$ along the entire segment ww'.

$$
\begin{aligned}
\mathcal{L}_2(w_t) &\le \begin{cases} (1-\alpha(t))\mathcal{L}_2(w) + \alpha(t)\mathcal{L}_2(w_{t_1}) & t \in [0, t_1] \\ (1-\beta(t))\mathcal{L}_2(w_{t_1}) + \beta(t)\mathcal{L}_2(w') & t \in [t_1, 1] \end{cases} \\
&< \begin{cases} (1-\alpha(t))\mathcal{L}_2(w) + \alpha(t)\left((1-t_1)\mathcal{L}_2(w) + t_1\mathcal{L}_2(w')\right) & t \in [0, t_1] \\ (1-\beta(t))\left((1-t_1)\mathcal{L}_2(w) + t_1\mathcal{L}_2(w')\right) + \beta(t)\mathcal{L}_2(w') & t \in [t_1, 1] \end{cases} \\
&= \begin{cases} (1-t)\mathcal{L}_2(w) + t\mathcal{L}_2(w') & t \in [0, t_1] \\ (1-t)\mathcal{L}_2(w) + t\mathcal{L}_2(w') & t \in [t_1, 1] \end{cases} \\
&= (1-t)\mathcal{L}_2(w) + t\mathcal{L}_2(w').
\end{aligned}
$$

The first inequality follows from the fact that $\mathcal{L}_2$ is convex along each sub-segment. The functions $\alpha : [0, t_1] \to [0, 1], \beta : [t_1, 1] \to [0, 1]$ are the reparameterization functions defined as $\alpha(t) = \frac{t}{t_1}, \beta(t) = \frac{t - t_1}{1 - t_1}$. The second inequality follows from the statement we proved about $\mathcal{L}_2(w_{t_1})$. The final equality follows from the definitions of α, β. Thus $\mathcal{L}_2$ is convex along the line segment ww' when $m = 1$.

Repeating the argument inductively gives that $\mathcal{L}$ is convex along ww' for any m. We have proven that the transitions between cells are strictly convex. When $\mathcal{L}_2$ restricted to each cell C is strongly convex, then the whole function $\mathcal{L}_2$ is strictly convex, otherwise $\mathcal{L}_2$ is convex. □

Lemma 68. *$\mathcal{L}_2$ is subdifferentiable everywhere. Let $w \in \mathbb{R}^d$. If $w \in \operatorname{Int} C$ for some $C \in \mathcal{H}$ and $w \neq 0$, then $\mathcal{L}_2$ is differentiable at w with*

$$\nabla \mathcal{L}_2(w) = X^\top(Xw - y) + \epsilon \|w\|_2 X^\top s + \epsilon \|Xw - y\|_1 \frac{w}{\|w\|_2} + \epsilon^2 n w. \tag{B.1}$$

If $w = 0$, then the subdifferential $\partial\mathcal{L}_2(0)$ is parameterized by replacing $\frac{w}{\|w\|_2}$ in Equation B.1 with any g such that $\|g\|_2 \leq 1$.

Otherwise $w \in \operatorname{Bd} C$, meaning that $w \in f$ for some $(d-k)$-dimensional face f of C. Let $\{i_1, \ldots, i_k\} \subset [n]$ be the k indices for which $w \in h_{i_j}$ ($x_{i_j}^\top w - y_{i_j} = 0$). The subdifferential $\partial\mathcal{L}_2(w)$ is non-empty and is parameterized by every setting of $s_{i_j} \in [-1, 1]$ in Equation B.1.

Proof. By Lemma 67, the epigraph of $\mathcal{L}_2$ is a convex set. The Separating Hyperplane Theorem implies the existence of a supporting hyperplane at every point $(w, \mathcal{L}_2(w))$. If $w \in \operatorname{Int} C$ for some cell C in the hyperplane arrangement, then $\mathcal{L}_2$ is differentiable at w and there is a single supporting hyperplane at $(w, \mathcal{L}_2(w))$. Otherwise, w is on the boundary ∂C of some C, and the existence of a supporting hyperplane implies the existence of a subgradient of $\mathcal{L}_2$ at w.

The gradient of $\mathcal{L}_2$, where defined, is

$$\nabla \mathcal{L}_2(w) = X^\top(Xw - y) + \epsilon \|w\|_2 X^\top s + \epsilon \|Xw - y\|_1 \frac{w}{\|w\|_2} + \epsilon^2 n w.$$

Suppose that $w = 0$. Let $v \in S^{d-1}$ be a unit vector and $\delta > 0$ sufficiently small. Then convexity of $\mathcal{L}_2|_{\operatorname{Int} C}$ for C with signature $s = -y$ (Lemma 66) and a standard limit argument gives

$$\begin{aligned}
\langle -X^\top y + \epsilon \|y\|_1 v, w' \rangle &= \lim_{\delta \to 0^+} \langle \nabla \mathcal{L}_2(\delta v), w' - \delta v \rangle \\
&\leq \lim_{\delta \to 0^+} \mathcal{L}_2(w') - \mathcal{L}_2(0) \\
&= \mathcal{L}_2(w') - \mathcal{L}_2(0)
\end{aligned}$$

where the inequality holds for all $\delta > 0$ by continuity. Thus v induces a subgradient at $w = 0$.

Let g be a vector such that $\|g\|_2 \leq 1$. g can be written as $g = (1-\alpha)v + \alpha(-v)$ for $\alpha = (1 - \|g\|_2)/2$ for subgradients $v, -v \in \partial\mathcal{L}_2(0)$. Since $\|g\|_2 \leq 1$, $\alpha \in [0, 1]$. So

$$\begin{aligned}
\langle -X^\top y + \epsilon \|y\|_1 g, w' \rangle &= \langle -X^\top y + \epsilon \|y\|_1 \left((1-\alpha)v + \alpha(-v)\right), w' \rangle \\
&= (1-\alpha)\left(\langle -X^\top y + \epsilon \|y\|_1 v, w' \rangle\right) + \alpha\left(\langle -X^\top y + \epsilon \|y\|_1 (-v), w' \rangle\right) \\
&\leq (1-\alpha)(\mathcal{L}_2(w') - \mathcal{L}_2(0)) + \alpha(\mathcal{L}_2(w') - \mathcal{L}_2(0)) \\
&= \mathcal{L}_2(w') - \mathcal{L}_2(0),
\end{aligned}$$

and so g induces a subgradient at $w = 0$ as well.

To find the subdifferential $\partial\mathcal{L}_2(w)$ for $w \neq 0$, we consider $w \in f$ for some $(d-k)$-dimensional facet f of C, and proceed by induction over k.

In the base case, $k = 1$. Since f is $(d-1)$-dimensional, there is only one tight hyerplane equation $x_i^\top w - y_i = 0$ at w. Let $h = \{w \in \mathbb{R}^d : x_i^\top w - y_i = 0\}$ denote the hyperplane and let h^+, h^- denote the halfspaces in which $\operatorname{sign}(x_i^\top w - y_i) = \pm 1$ respectively. The limit of the gradient as approach w by a sequence in h^+ is $\nabla\mathcal{L}_2(w)$ where $s_i = 1$; similarly approaching w by a sequence in h^- gives $\nabla\mathcal{L}_2(w)$ where $s_i = -1$. These vectors define two supporting hyperplanes of the epigraph at $(w, \mathcal{L}_2(w))$.

Note that only $\epsilon\|w\|_2 X^\top s$ and $\epsilon\|Xw - y\|_1 \frac{w}{\|w\|_2}$ in $\nabla\mathcal{L}_2$ depend upon s, and when $x_i^\top w - y_i = 0$, $\|Xw - y\|_1$ is identical regardless of the setting of s_i, so we need only consider $\epsilon\|w\|_2 X^\top s$. Let $s_i \in [-1, 1]$, then

$$\begin{aligned}
\epsilon\|w\|_2\langle X^\top s, w' - w\rangle &= \epsilon\|w\|_2\langle\left(\sum_{j\neq i} s_j x_j\right) + s_i x_i, w' - w\rangle \\
&= \epsilon\|w\|_2\left(\langle\sum_{j\neq i} s_j x_j, w' - w\rangle + \langle s_i x_i, w' - w\rangle\right) \\
&= \epsilon\|w\|_2\left(\langle\sum_{j\neq i} s_j x_j, w' - w\rangle + (1-\alpha)\langle -x_i, w' - w\rangle + \alpha\langle x_i, w' - w\rangle\right)
\end{aligned}$$

where $\alpha = \frac{1+s_i}{2}$. Then we can express $\nabla\mathcal{L}_2(w)|_{s_i}$, where $s_i \in [-1, 1]$, as a convex combination of the terms $\mathcal{L}_2(w)|_{s_i=-1}, \mathcal{L}_2(w)|_{s_i=1}$.

$$\begin{aligned}
\langle\nabla\mathcal{L}_2(w)|_{s_i}, w' - w\rangle &= \langle X^\top(Xw - y) + \epsilon\|w\|_2 X^\top s + \epsilon\|Xw - y\|_1\frac{w}{\|w\|_2} + \epsilon^2 n w, w' - w\rangle \\
&= \langle X^\top(Xw - y) + \epsilon\|Xw - y\|_1\frac{w}{\|w\|_2} + \epsilon^2 n w, w' - w\rangle \\
&\quad + \epsilon\|w\|_2\left(\langle\sum_{j\neq i} s_j x_j, w' - w\rangle + (1-\alpha)\langle -x_i, w' - w\rangle + \alpha\langle x_i, w' - w\rangle\right) \\
&= (1-\alpha)\langle\nabla\mathcal{L}_2(w)|_{s_i=-1}, w' - w\rangle + \alpha\langle\nabla\mathcal{L}_2(w)|_{s_i=1}, w' - w\rangle \\
&\leq (1-\alpha)\left(\mathcal{L}_2(w') - \mathcal{L}_2(w)\right) + \alpha\left(\mathcal{L}_2(w') - \mathcal{L}_2(w)\right) \\
&= \mathcal{L}_2(w') - \mathcal{L}_2(w)
\end{aligned}$$

where the inequality follows from the fact that $\nabla\mathcal{L}_2(w)|_{s_i=-1}, \nabla\mathcal{L}_2(w)|_{s_i=1}$ are subgradients. Thus $\mathcal{L}_2(w)|_{s_i}$ is a subgradient for any $s_i \in [-1, 1]$ at w.

Now suppose that $w \in f$ is a $(d-k)$-dimensional facet and the statement holds for all $1 \leq j < k$. Let $\{i_1, \ldots, i_k\}$ index the hyperplane equations $x_{i_j}^\top w - y_{i_j} = 0$ at w. Consider the subset of hyperplane equations $\{i_1, \ldots, i_{k-1}\}$ along which subgradients exist for any setting of $s_{i_j} \in [-1, 1]$ by the inductive hypothesis. An identical limit argument as above implies the existence of two subgradients at w with $s_{i_k} = \pm 1$. Then an identical calculation to those above imply that $\nabla\mathcal{L}_2(w)|_{s_{i_k}}$ is a subgradient for any $s_{i_k} \in [-1, 1]$. Thus at $w \in f$ there exists a subdifferential parameterized by $s_{i_j} \in [-1, 1]$ for every $1 \leq j \leq k$. □

Appendix C

Corrections

[93] derive the minimum norm solution using the kernel trick. The optimal solutions $w_{\text{SDG}} = X^\top \alpha$ where $\alpha = K^{-1}y$ for $K = XX^\top$. They compute

$$K_{ij} = \begin{cases} 4 & \text{if } i = j \text{ and } y_i = 1 \\ 8 & \text{if } i = j \text{ and } y_i = -1 \\ 3 & \text{if } i \neq j \text{ and } y_i y_j = 1 \\ 1 & \text{if } i \neq j \text{ and } y_i y_j = -1 \end{cases}$$

and positing, correctly, that $\alpha_i = \alpha_+$ if $y_i = 1$ and $\alpha_i = \alpha_-$ if $y_i = -1$ they derive the system of equations

$$\begin{aligned} (3n_+ + 1)\alpha_+ + n_-\alpha_- &= 1 \\ n_+\alpha_+ + (3n_- + 3)\alpha_- &= -1 \end{aligned}$$

which gives

$$\alpha_+ = \frac{4n_- + 3}{9n_+ + 3n_- + 8n_+n_- + 5}, \quad \alpha_- = -\frac{4n_+ + 1}{9n_+ + 3n_- + 8n_+n_- + 5}.$$

[93] mistakenly dropped the negative in α_-. Unfortunately there is an additional minor mistake in the linear system. The system is derived by computing

$$\begin{aligned} (K\alpha)_i &= \begin{cases} 4\alpha_i + \sum_{j\in\mathcal{P}-i} 3\alpha_j + \sum_{j\in\mathcal{N}} \alpha_j & \text{if } y_i = 1 \\ 8\alpha_i + \sum_{j\in\mathcal{P}} \alpha_j + 3\sum_{j\in\mathcal{N}-i} \alpha_j & \text{if } y_j = 1 \end{cases} \\ &= \begin{cases} \alpha_i + \sum_{j\in\mathcal{P}} 3\alpha_j + \sum_{j\in\mathcal{N}} \alpha_j & \text{if } y_i = 1 \\ 5\alpha_i + \sum_{j\in\mathcal{P}} \alpha_j + 3\sum_{j\in\mathcal{N}} \alpha_j & \text{if } y_j = 1 \end{cases}. \end{aligned}$$

Subtracting equations we reach the conclusion that $\alpha_i = \alpha_+$ if $y_i = 1$ and $\alpha_i = \alpha_-$ if $y_i = -1$. Then it's clear that there are really only two equations in this system

$$\begin{aligned} (3n_+ + 1)\alpha_+ + n_-\alpha_- &= 1 \\ n_+\alpha_+ + (3n_- + 5)\alpha_- &= -1 \end{aligned}$$

which gives

$$\alpha_{+} = \frac{4n_{-} + 5}{15n_{+} + 3n_{-} + 8n_{+}n_{-} + 5}, \quad \alpha_{-} = -\frac{4n_{+} + 1}{15n_{+} + 3n_{-} + 8n_{+}n_{-} + 5}.$$