

**TIGHTENING AND TIGHT HULLS: MAXIMIZING
CONVEXITY SUBJECT TO SET-THEORETIC
CONSTRAINTS**

A Dissertation

Presented to

The Academic Faculty

By

Jason D. Williams

In Partial Fulfillment

Of the Requirements for the Degree

Doctor of Philosophy in Computer Science

Georgia Institute of Technology

August 2010

**TIGHTENING AND TIGHT HULLS: MAXIMIZING
CONVEXITY SUBJECT TO SET-THEORETIC
CONSTRAINTS**

Approved by:

Date Approved:

The art of doing mathematics consists in finding that special case which contains all the germs of generality. — D. Hilbert, in *Mathematical Maxims and Minims* by N. Rose.

Dedicated to...

ACKNOWLEDGEMENTS

First, I must thank...

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	v
LIST OF FIGURES	x
NOTATION	xiv
SUMMARY	xvi
PART 1. OVERVIEW	1
CHAPTER 1. INTRODUCTION	2
CHAPTER 2. DOCUMENT STRUCTURE	6
2.1. Division into parts	6
2.2. Part I	7
2.3. Part II	8
2.4. Part III	10
2.5. Part IV	11
CHAPTER 3. TIGHT HULLS OVERVIEW	12
3.1. Tight hulls generalize convex hulls	12
3.2. Comparison of hull definitions	15
3.3. Hulls of revolution	30
3.4. Symmetry with respect to set complement	36
3.5. The tablecloth problem	38
CHAPTER 4. TIGHTENING OVERVIEW	45
4.1. Tightening bounds curvature	45

4.2.	Continuously tightening polygons	49
4.3.	Tightening simplifies normal fields	53
CHAPTER 5. MEDIAL COVER OVERVIEW		58
CHAPTER 6. IMPLEMENTATION OVERVIEW		65
CHAPTER 7. APPLICATIONS		74
7.1.	Convergent boundary estimation	74
7.2.	Shape design	79
7.3.	Blending	80
7.4.	Normal field simplification	81
7.5.	Further applications	81
PART 2. BACKGROUND		84
CHAPTER 8. CONVEX STRUCTURES AND CLOSURE STRUCTURES		85
8.1.	Overview	85
8.2.	Ball notation	86
8.3.	Interior, closure, opening, and closing; duality	87
8.4.	Dilation, erosion, and Minkowski sum and difference	89
8.5.	Convex hulls, relative convex hulls, and affine hulls	94
8.6.	Closure structures	96
8.7.	Convex structures	99
8.8.	Regularity and symmetry	102
CHAPTER 9. THE MEDIAL AXIS		108
9.1.	Overview	108
9.2.	Distance and regularity transforms	108
9.3.	Maximal balls and the medial axis	113
9.4.	Bloom, contact set, and radius function	114
9.5.	Comb and alpha shape	116

9.6. Piecewise real analytic boundary assumption	117
CHAPTER 10. NORMALS AND CURVATURES	121
10.1. Overview	121
10.2. Hausdorff measure, geodesic balls, and dimensionality	122
10.3. Normals for arbitrary subsets of Euclidean space	127
10.4. Slack and Gaussian curvature	138
10.5. Mean and sectional curvature	144
10.6. Lipschitz continuity and smooth offset interval	147
CHAPTER 11. TIGHT EMBEDDINGS	151
11.1. Overview	151
11.2. Smooth and polyhedral tightness	152
11.3. Minimizing height function maxima in \mathbb{R}^3 and \mathbb{R}^d	164
11.4. Height function critical points	170
CHAPTER 12. COVERS AND HULLS	177
12.1. Hulls and relative hulls	177
12.2. Support and convexity normals	190
12.3. Topological and geometric distinguishability	197
12.4. Covers locally indistinguishible from hulls	197
PART 3. TIGHT HULLS AND TIGHTENING	200
CHAPTER 13. TIGHT HULLS	201
CHAPTER 14. TIGHTENING	202
CHAPTER 15. THE MEDIAL COVER	203
CHAPTER 16. IMPLEMENTATION	204
PART 4. PRIOR ART, FUTURE WORK, CONTRIBUTIONS, AND CONCLUSION	205

CHAPTER 17. PRIOR ART: FAIRING, BLURRING, AND BLENDING	206
CHAPTER 18. FUTURE WORK: CONSTRUCTING 3D TIGHT HULLS	207
CHAPTER 19. CONTRIBUTIONS	208
CHAPTER 20. CONCLUSION	209
REFERENCES	210

LIST OF FIGURES

3.1	A convex hull of points in the plane.	12
3.2	A hollow cylinder swept by rotating thin rectangles.	13
3.3	The convex hull of a hollow cylinder.	13
3.4	A separating surface with fluid-like behavior.	14
3.5	A separating surface like a membrane in tension.	14
3.6	Definition of a convex set.	16
3.7	Definition of a convex hull.	16
3.8	Convex hull construction by filtering subsets of Euclidean space.	17
3.9	Definition of relative convexity.	18
3.10	The convex hull of one simple polygon relative to another.	18
3.11	Embeddings of the torus that are and are not tight.	20
3.12	Embeddings of the sphere that are and are not tight.	21
3.13	Integral Gaussian curvature at a triangle mesh vertex.	21
3.14	Gaussian curvature from geodesic circles.	22
3.15	The total absolute curvature of a sphere and a cube are equal.	23
3.16	Sets that have and do not have the two-piece property.	24
3.17	A problem with defining tight hulls using the two-piece property.	24
3.18	Uncountably many sets may minimize constrained slack.	25
3.19	The support of a normal to a candidate tight hull.	26
3.20	Unsupported slack for different candidate tight hulls.	27
3.21	Convex hulls have supported slack.	28

3.22	A singular geometric configuration yielding a nonunique tight hull.	29
3.23	An intersection of tight sets that is not tight.	30
3.24	The input sets R and G are symmetric about the same line.	31
3.25	The tight hull and convex hull of R relative to G are identical.	32
3.26	The sets R' and G' swept by rotating R and G .	32
3.27	The set R' consists of hollow and solid cylinders.	33
3.28	The set G' consists of a hollow cylinder and the complement of a cylinder.	34
3.29	The convex hull of R' relative to G' .	34
3.30	The tight hull of R' relative to G' .	35
3.31	Cross sections of the tight hull and convex hull of R' relative to G' .	35
3.32	The tight hull and convex hull of G' relative to R' .	37
3.33	The complements of the tight hull and convex hull of R' relative to G' .	37
3.34	The convex hull of four equal-radius balls.	38
3.35	The set R consists of four red balls, while G is one green ball.	39
3.36	The convex hull of R relative to G .	39
3.37	The tight hull of R relative to G .	40
3.38	Elliptic cylinders in the tight hull.	40
3.39	An incorrect tight hull of R relative to G with circular cylinders.	41
3.40	A spherical cap in the incorrect hull is larger than in the tight hull.	42
3.41	Singular arcs on the incorrect hull have excess normal variation.	43
3.42	The tight hull of a torus relative to a ball.	44
4.1	Hull boundaries press against locally convex constraint boundaries.	45
4.2	Planar hull boundaries are stable under constrained curvature flow.	46
4.3	The normals to a two-dimensional hull are supported.	47
4.4	The morphological opening and closing of a polygon.	48

4.5	The tightening of a polygon.	49
4.6	A set's mortar produced by rolling balls inside and outside of it.	50
4.7	The curvature of a hull's boundary is bounded by its constraints.	51
4.8	The normals to a tightening may be continuous but nondifferentiable.	51
4.9	A tightened polygon's boundary is an sequence of segments and arcs.	52
4.10	Arcs in a polygon's tightening distribute the normal change at its vertices.	52
4.11	Increasing the tightening radius slides arc endpoints together.	53
4.12	Changing the tightening radius can change the result's topology.	54
4.13	The normal field to a polygon and its tightening.	55
4.14	The Gauss offset of a nonconvex polygon.	55
4.15	The normal count from ray intersections with the Gauss offset.	56
4.16	Several paths connecting edge endpoint have identical slack.	57
4.17	Unsupported edges trim slack from incident arcs.	57
5.1	An algorithm for the relative convex hull of simple polygons.	59
5.2	Tight covers locally behave like tight hulls.	60
5.3	There may be an exponential number of tight covers.	60
5.4	Tight cover boundaries wind around holes in G^c .	61
5.5	Points added to R and G to define the medial cover.	62
5.6	The medial cover boundary is isotopic to loops equidistant from R and G .	63
5.7	The ordering of planar tight covers by set inclusion.	64
6.1	Red, green, and mixed triangles from the triangulation of $(R \cup G)^c$.	65
6.2	Mixed triangles form topological annuli.	66
6.3	The constrained Delaunay cover.	67
6.4	The constrained Delaunay cover depends on sampling.	68

6.5	Constrained Delaunay triangulation converges at high sampling rates.	69
6.6	Large triangle circumcenters converge to medial axis bifurcation points.	69
6.7	Insertion of bifurcation disk contact points as polygon vertices.	70
6.8	Bifurcation disk contact points make a cover invariant to sample insertion.	70
6.9	A medial cover algorithm can compute the convex hull.	71
6.10	Opening with an n -sided polygon has an output size linear in n .	72
6.11	Hausdorff error due to opening with a regular polygon rather than a disk.	72
6.12	Discrete analogs of bounded curvature in polygonal tightening.	73
7.1	Morphologically regular and irregular sets.	74
7.2	Red and green sets from samples.	75
7.3	Relative convex hull boundary measure converges with sampling density.	76
7.4	Peano-Jordan measure.	77
7.5	Midpoint reconstruction is nonconvergent.	78
7.6	Designs incorporating developable surfaces.	79
7.7	Developable surfaces interpolating boundary patches.	80
7.8	Small-radius tightening provides unique normals.	81
7.9	Tightening progressively simplifies normal fields.	82
7.10	Repairing polygonal data with tight hulls.	82
8.1	Closure and interior operators do not always produce manifold boundaries.	104
10.1	A geodesic distance singularity separating points closest to the same site.	125
10.2	Morphological closure does not provide normals at concave singularities.	130
10.3	Normals defined by ϵ -closure of an arc of radius ϵ point to the arc's center.	131
10.4	A set with arbitrarily high slack over the neighborhood of a point.	135
10.5	Normal count is the sum over normal patches of patch normal count.	137

NOTATION

$X \setminus Y$	set difference; the elements of X not contained in Y
X^c	the complement of set X ; for universal set U , $X^c := U \setminus X$
$\wp(X)$	powerset of X : $\wp(X) := \{Y \mid Y \subseteq X\}$
$\#X$	cardinality of X
$\kappa(X)$	the set of connected components of X
\mathbb{R}	the set of real numbers
$\mathbb{R}^<, \mathbb{R}^{\leq}$	the set of real numbers less than/less than or equal to zero
$\mathbb{R}^>, \mathbb{R}^{\geq}$	the set of real numbers greater than/greater than or equal to zero
\mathbb{R}^d	d -dimensional Euclidean space
\mathbb{S}^d	d -dimensional unit sphere bounding the unit ball in \mathbb{R}^{d+1}
\mathbb{Q}	rational numbers; $\mathbb{Q}^d \subseteq \mathbb{R}^d$ is the set of points with rational coordinates
$ \cdot $	absolute value
$\ \cdot\ $	Euclidean vector norm; for $v \in \mathbb{R}^d$, $\ v\ := \left(\sum_{i=1}^d v_i^2\right)^{\frac{1}{2}}$
\bigwedge	infimum; greatest lower bound
\bigvee	supremum; least upper bound
$d(\cdot, \cdot)$	Euclidean distance function for points and point sets
$cp^S(p)$	set of points from $S \subseteq \mathbb{R}^d$ closest to point $p \in \mathbb{R}^d$
$g^S(\cdot, \cdot)$	geodesic distance over $S \subseteq \mathbb{R}^d$
$g_r^S(p)$	geodesic ball of radius $r \in \mathbb{R}^>$ in $S \subseteq \mathbb{R}^d$ centered on $p \in S$
$O(f(x))$	set of functions asymptotically bounded from above by a multiple of $f(x)$
$\Omega(f(x))$	set of functions asymptotically bounded from below by a multiple of $f(x)$
$\Theta(f(x))$	functions bounded above and below; $\Theta(f(x)) := O(f(x)) \cap \Omega(f(x))$

$S + v$	set $S \subseteq \mathbb{R}^d$ translated by vector $v \in \mathbb{R}^d$
\overline{pq}	the closed line segment connecting points $p, q \in \mathbb{R}^d$
\overleftrightarrow{pq}	the line through distinct points $p, q \in \mathbb{R}^d$
\overrightarrow{pv}	the ray with endpoint $p \in \mathbb{R}^d$ and direction $\frac{v}{\ v\ }$ for vector $v \in \mathbb{R}^d$
$B_r(p)$	d -dimensional open ball of radius $r \in \mathbb{R}^>$ centered on $p \in \mathbb{R}^d$
$\rho(b)$	radius of ball b
$\pi(b)$	center of ball b
\mathbb{B}	set of all open balls; $\mathbb{B} := \{B_r(p) \mid (r \in \mathbb{R}^>) \wedge (p \in \mathbb{R}^d)\}$
\mathbb{B}^S	open balls contained in $S \subseteq \mathbb{R}^d$; $\mathbb{B}^S := \{b \in \mathbb{B} \mid b \subseteq S\}$
\mathbb{B}_r	open balls of radius $r \in \mathbb{R}^>$; $\mathbb{B}_r := \{b \in \mathbb{B} \mid \rho(b) = r\}$
$\mathbb{B}_{>r}, \mathbb{B}_{\geq r}$	open balls with radius greater than/greater than or equal to r
$\mathbb{B}_{<r}, \mathbb{B}_{\leq r}$	open balls with radius less than/less than or equal to r
$\mathbb{B}_r^S, \mathbb{B}_{\leq r}^S, \dots$	open balls in S with radius r , radius less than or equal to r , etc.
S°	interior of $S \subseteq \mathbb{R}^d$; $S^\circ := \bigcup \mathbb{B}^S$
S^-	closure of $S \subseteq \mathbb{R}^d$; $S^- := S^{coc}$
∂S	boundary of $S \subseteq \mathbb{R}^d$; $\partial S := S^- \cap S^{c-}$
$S \circ_r$	morphological opening of $S \subseteq \mathbb{R}^d$ with radius $r \in \mathbb{R}^>$; $S \circ_r := \bigcup \mathbb{B}_r^S$
$S \bullet_r$	morphological closure of $S \subseteq \mathbb{R}^d$ with radius $r \in \mathbb{R}^>$; $S \bullet_r := S^c \circ_r^c$
$M_r S$	mortar of $S \subseteq \mathbb{R}^d$ for radius $r \in \mathbb{R}^>$; $M_r S := S \bullet_r \cap S^c \bullet_r$

SUMMARY

We introduce tight hulls, which generalize convex hulls. While the convex hull of a subset R of Euclidean space is the smallest convex set containing R , a tight hull of R relative to a second set G contains R , excludes G , and maximizes convexity by minimizing a generalization of total absolute curvature we call slack. Out of all sets with those properties, tight hulls additionally minimize a quantity we call unsupported slack. This makes the boundary of a tight hull behave like a membrane in tension, stretched over R and G . Tight hulls are closely related to relative convex hulls, but possess a symmetry with respect to set complement that relative convex hulls lack. Our hope is that this symmetry enhances three-dimensional boundary estimation from volumetric samples and provides a new way to formally specify a useful class of solid models. From tight hulls we obtain r -tightening: the r -tightening of a set is the tight hull of its morphological opening with an open ball of radius r relative to the opening of its complement. We extend our previous results on the curvature-limiting properties of tightening by showing that continuously increasing the radius r used for tightening produces a piecewise continuous simplification of the tightened set's normal field. We relax the definition of a tight hull to obtain tight covers, which locally behave like tight hulls but have more diverse topologies. We identify the medial cover in two dimensions as a tight cover symmetric with respect to set complement, and we show that for polygonal inputs it can be robustly computed in $O(n \log n)$ time using standard algorithms. We approximate the opening of a polygon with a disk by replacing the disk with a regular polygon, then apply our medial cover algorithm to the polygonal opening to perform robust, approximate tightening.

PART 1

OVERVIEW

CHAPTER 1

INTRODUCTION

Problem and contributions. In this dissertation, we address the problem of maximizing convexity subject to set-theoretic constraints. We present five contributions: tight hulls, tightening, generalized normals and curvatures, a framework for covers and hulls, and robust $O(n \log n)$ algorithms for constructing tight covers and approximate tightenings from two-dimensional polygonal inputs.

Tight hulls. Tight hulls and tightenings are our main contributions. Tight hulls generalize convex hulls; like a convex hull, a tight hull is a maximally convex set that contains a given set. Unlike a convex hull, a tight hull is also constrained to exclude a second set. The boundary of a tight hull separates the interiors of the sets it includes and excludes, and it stretches over those sets like a membrane in tension. We anticipate applications of tight hulls to solid model design and convergent boundary estimation.

Tightening. The tightening of a set with radius r is a tight hull that contains the balls of radius r lying in the set's interior and excludes the balls of radius r lying in its exterior. Interpreting r as a scale parameter, a set's deformation as it is tightened is piecewise continuous as a function of scale. We show that progressively increasing the value of r simplifies a set's normal field in two ways. First, the total variation of a set's normal field monotonically decreases as a function of r over intervals where the tightening deforms continuously. Second, tightening a two-dimensional set bounds the rate of change in normal direction over its boundary as a function of arclength. We anticipate applications of tightening to solid blending and normal field analysis.

Generalized normals and curvatures. To support our work on tight hulls and tightening, we provide generalized definitions of normal fields and normal variation, as well

as a framework for comparing different hull-like sets. In our definition of generalized normal fields, we map each point on the boundary of an arbitrary subset of Euclidean space to a nonempty set of normals. We then define mean curvature, Gaussian curvature, sectional curvature, normal field Lipschitz continuity, and the smooth offset interval in ways that make them applicable to sets whose boundaries include points with nonunique normals. We ensure that our definitions generalize smooth definitions, and we preserve important results from classical differential geometry, such as the Gauss-Bonnet Theorem and the Chern-Lashof Inequalities. Our definitions provide normal fields for polyhedra, enabling us to compute integral curvatures over their edges and vertices.

Covers and hulls. We define hulls using the hull property. In essence, the hull property guarantees that a plane cannot shave a connected component from a hull without intersecting the hull's constraints. We show that all hulls have the same local geometric properties in two dimensions, but there are uncountably many different hull geometries in three and higher dimensions. We further show that hulls with the same local geometry may have different topologies, and we define covers to separate a hull's topology from its geometry. We identify the medial cover as a cover that is easy to compute in two dimensions using the medial axis, and we show that the medial cover shares some of the tight hull's most desirable properties.

Implementation. We describe simple, efficient, and robust algorithms for computing the medial cover and a variant of tightening called medial tightening in two dimensions for input constraints with polygonal boundaries. Given robust implementations of algorithms for the medial axis, constrained Delaunay triangulation, and path planning, robustly computing the medial cover in worst-case $O(n \log n)$ time requires little additional development. Computing a polygonal form of medial r -tightening additionally requires an algorithm for morphologically opening an arbitrary polygon with a regular polygon approximating a disk of radius r . We use the output of the opening algorithm as input to the medial cover algorithm.

Applications. A few promising applications require us to compute tight hulls in three dimensions. We know of no efficient algorithm for computing three-dimensional tight hulls, but we have reasons to believe the problem is tractable. We conjecture that the portion of a tight boundary that is disjoint from the boundary of the input constraints is developable, and we further conjecture that tight hulls of polyhedra are polyhedral. In contrast to other fairing problems that generate surfaces that must be sampled densely to be represented accurately, a nearly exact tight hull representation promises in practice to be comparable in size to the input.

Assuming we can compute tight hulls in three dimensions, they may prove valuable as a convergent boundary estimators. If the surface area of a tight hull, like the surface area of a relative convex hull, converges to the surface area of a rasterized solid as the distance between the samples used for rasterization goes to zero, tight hulls can numerically solve definite integrals that are insoluble in closed form [34]. If the normal field of a tight hull converges to the normal field of a rasterized solid, tight hulls may also prove useful in medical imaging. Provided with a segmented medical image that labels each point on a lattice as in or out of a particular anatomical structure, a tight hull that includes the “in” samples and excludes the “out” samples can produce an image of the the structure whose shading is guaranteed to be accurate because the hull’s normal field is accurate.

If the surface patches introduced by tight hulls are developable, tight hulls may find use manufacturing. Developable surfaces can be produced by cutting, bending, and folding flexible sheets. These processes are often more affordable than other manufacturing techniques, such as milling and casting. The increasing availability of programmable cutting tools also suggests the possibility of producing objects composed of developable surfaces without the need for large capital investments.

Audience and context. We draw on the fields of solid modeling, computer graphics, computational geometry, and differential geometry, with significant applications to biomedical imaging. We approach convexity from the perspectives of both convex

structures [63] and tight embeddings [3]. Building on the theory of tight embeddings, we devise definitions of normals and normal variation applicable to polyhedra, rederiving results from the field of discrete differential geometry [6]. We employ the medial axis and related constructs such as constrained Delaunay triangulation, so we include a review of medial axis theory. This dissertation may be useful to readers interested in any of these fields.

CHAPTER 2

DOCUMENT STRUCTURE

2.1. Division into parts

This dissertation is divided into four parts:

Part I. In Part I, we provide an overview of our main technical contributions. We define tight hulls and tightening, and we describe their most significant properties. We introduce the medial cover and detail how we can use it to construct variants of tight hulls and tightening in two dimensions given polygonal inputs.

Part I documents the core of this dissertation. It includes precise claims and definitions, but it is designed to be easy to read. We minimize our use of notation and mathematical English, deferring proofs and technical refinement to Parts II and III. We also densely illustrate the text.

Part II. In Part II, we formalize the material we need to rigorously present tightening and tight hulls. We include a combination of prior art and original work. Our chapters on convex structures, the medial axis, and tight embeddings focus on prior art, while our chapters on normals and covers and hulls are largely original. Part II introduces much of our notation, and it describes any nonstandard usage we make of standard terms.

Part III. In Part III, we formally define tight hulls and tightening, and we prove theorems concerning some of their most significant properties. We also define the medial cover in terms of the medial axis. We detail the robust and efficient construction of two-dimensional medial covers and medial tightenings, and we prove that our construction is correct.

Part IV. Part IV outlines a context for tight hulls and tightening. We begin by comparing our work to prior art that shares similar goals but addresses them through

different means. Adopting the perspective that some of the most promising applications for maximally convex sets are in three dimensions, we describe our current understanding of what computing three-dimensional tight hulls entails. We summarize our contributions, then conclude with a discussion of issues raised by our work.

The next four sections briefly describe the purpose and contents of each chapter.

2.2. Part I

Introduction. We summarize our technical contributions and prospective applications, indicate the relationship of our work to established disciplines, and characterize our target audience.

Document structure. We describe this document's organization.

Tight hulls overview. We introduce the tight hull using elements of the theory of tight embeddings, comparing it to the convex hull and relative convex hull. Focusing on tight hull of rotationally symmetric sets, we explain our claim that the tight hull is symmetric with respect to set complement while the relative convex hull is not. We conclude by presenting a more complex tight hull that is not rotationally symmetric, establishing that three-dimensional tight hulls defined by smooth constraints are not necessarily smooth.

Tightening overview. We define r -tightening in terms of tight hulls and morphological opening, and we illustrate how in two dimensions the r -tightening of a set has a curvature between $-\frac{1}{r}$ and $\frac{1}{r}$. We describe the deformation of the r -tightening of a polygon as we continuously increase the value of r , claiming that the total variation in the normal field of an r -tightening monotonically decreases over intervals of r where the tightening varies continuously. Combining this result with the curvature-limiting properties of tightening, we conclude that tightening simplifies normal fields.

Medial cover overview. We introduce two-dimensional tight covers, which have the same local geometric properties as the two-dimensional tight hull but may not share the tight hull's topology. As a special case, we identify the medial cover, whose

bounding loops can be continuously deformed to the points equidistant from the set constrained to lie inside the cover and the set constrained to lie outside of it. We show the medial cover, like the tight hull, is symmetric with respect to set complement, while the relative convex hull is not.

Implementation overview. We explain how to construct a two-dimensional tight cover from the triangulation of the gap between two polygonal sets, and we show that the cover topology depends on the triangulation. We eliminate the variability due to different triangulations by adding to the input points obtained from the medial axis of the gap between the constraints. Constructing the constrained Delaunay triangulation of the augmented input, we obtain the medial cover. We obtain an approximate form of tightening by opening the input with a regular polygon, and we characterize the resulting error.

Applications. We describe the application of tight hulls to convergent boundary measure estimation, speculating that the tight hull may have a convergent normal field as well. Assuming that the portion of a three-dimensional tight hull's boundary disjoint from its constraints is developable, we describe applications of tight hulls to the design of shapes that can be manufactured from flexible sheets. We propose that tightening can function as a symmetric solid blend, and we suggest that it may prove useful in the multiscale analysis and simplification of normal fields.

2.3. Part II

Convex structures and closure structures. We begin by defining topological and morphological operators in terms of open balls. We then define the Euclidean convex hull, the relative convex hull, the geodesic hull, and the affine hull, highlighting the similarities between them. We use convex structures and closure structures [63] to capture properties shared by various operators on the lattice of subsets of Euclidean space, concluding with material relevant to the relationship between duality and symmetry with respect to set complement.

The medial axis. We introduce the regularity transform, which quantifies the thickness of a set, and relate it to the distance transform. We then define maximal balls and the medial axis, as well as terminology pertaining to the medial axis. We introduce the comb of a subset of the medial axis and use it to define a continuous version of the α -shape [19, 20]. We conclude by explaining why a set whose boundary can be expressed as the union of a finite number of real analytic pieces has a medial axis with a finite graph structure [15].

Normals and curvatures. We assign each point on the boundary of an arbitrary subset of Euclidean space a nonempty set of normals. We show that our definition generalizes definitions of normal fields for smooth sets, producing normals at singularities comparable to those we might obtain by slightly blending the singular surface. We define integral and pointwise curvatures so that they exhibit meaningful behavior at singularities, and we explain their application to polyhedra. We conclude by defining the Lipschitz continuity of a normal field and introducing the smooth offset interval.

Tight embeddings. Beginning with a review of curvature for smooth surfaces using the shape operator, we define tight surfaces as minimizers of total absolute curvature. Providing formulas for the integral Gaussian curvature and absolute Gaussian curvature at a polyhedron's vertices, we similarly define tight polyhedra. We unify tight surfaces and tight polyhedra in low dimensions using the two-piece property, then relate the two-piece property to the minimization of height function critical points. We conclude by showing that minimizing height function critical points is nearly identical to minimizing our convexity functional.

Covers and hulls. We introduce hulls, which formalize the idea of wrapping the boundary of a set around a second set in its interior. We then introduce relative hulls, whose boundaries wrap around two sets, separating them. We show that relative convex hulls and tight hulls are relative hulls. We then define the support a hull receives from its constraints, proving that a convex hull is fully supported by the set it contains. We define when two hulls can be distinguished geometrically or topologically,

and we show that two-dimensional hulls are geometrically indistinguishable. Finally, we define covers to group hulls into geometrically indistinguishable classes.

2.4. Part III

Tight hulls. We define tight hulls as sets that minimize first slack and then unsupported slack subject to volumetric constraints. We explain the grounds for our conjecture that the unsupported boundary of a tight hull is developable, and we prove several claims made in our overview of tight hulls (Chapter 3.)

Tightening. Focusing on two-dimensions, we prove that a set's r -tightenings are piecewise continuous as a function of r . We further show that the tightenings' slack monotonically decreases over intervals where the tightening is continuous. We associate this decrease in slack with the unfolding of a set we call the ϵ -Gauss offset. We prove that the normal field to an r -tightening is Lipschitz continuous with a constant of $\frac{1}{r}$, and we show that the circular arcs introduced by tightening linearly interpolate the normals at their endpoints.

The medial cover. We define the medial cover in two dimensions using the medial axis of the gap between the sets constrained to lie inside and outside of the cover. We prove that we can deform the boundary of the medial cover to the set of points equidistant from the contained and excluded constraints while keeping them respectively contained and excluded throughout the deformation. We demonstrate that, unlike the relative convex hull, the medial cover does not have a fixed position among all covers when they are ordered by set inclusion. Instead, it groups connected components of the constraints by proximity.

Implementation. We define solid and mixed triangles in the triangulation of the gap between polygonal constraints, then show that mixed triangles connect across mixed edges to form annuli. We prove that if we insert contact points of maximal disks centered on medial axis bifurcation points into the input, the constrained Delaunay triangulation of the augmented input yields annuli that can be contracted to the

bounding loops of the medial cover. We close by proving that approximate tightenings constructed with regular polygons possess discrete analogs of the properties of tightenings constructed with disks.

2.5. Part IV

Prior art: Fairing, blurring, and blending. We place tight hulls in the context of surface optimization by comparing slack to existing surface functionals, focusing on minimal surfaces [26, 68] and minimum variation of curvature surfaces [42, 43]. We then compare tightening to lowpass surface filtering through the modulation of Laplace-Beltrami eigenfunction amplitudes [16, 40], proving that lowpass filters do not guarantee curvature bounds. Finally, we compare tightening to solid blends involving rolling balls [53, 52], emphasizing the distinctiveness of tightening’s symmetry with respect to set complement.

Future work: Constructing 3D tight hulls. We propose reasons why the efficient construction of three-dimensional relative convex hulls remains an open problem despite demand for convergent boundary reconstruction from volumetric samples. We show that the tight hull, like the relative convex hull, may include $O(n^2)$ vertices that are not part of the input. We describe an inefficient polynomial time algorithm for constructing tight hulls, and we demonstrate a tractable iterative technique for the construction of approximate relative convex hulls.

Contributions. We review our contributions and their significance.

Conclusion. We discuss three issues raised our work. The first concerns the roles of slack and unsupported slack in problems that involve constrained and approximate convexity. The second concerns the design of symmetric versions of closure operators. The third concerns the construction of set-theoretic deformations from locally-applied closure operators.

In the next chapter, we begin an overview of our main contributions.

TIGHT HULLS OVERVIEW

3.1. Tight hulls generalize convex hulls

We define tight hulls as a generalization of Euclidean convex hulls. The convex hull of a set $S \subseteq \mathbb{R}^d$ can be characterized as the smallest convex set containing S (Figure 3.1.) A tight hull is also a maximally convex set of minimum size. However, while the convex hull of S is defined so that it contains S , the tight hull of $R \subseteq \mathbb{R}^d$ relative to a set $G \subseteq \mathbb{R}^d$ disjoint from R is defined to contain R and exclude G . If the interior of G intersects the convex hull of R , the tight hull cannot be convex, requiring us to formalize how a nonconvex set can be maximally convex given a set of constraints. How we define convexity then determines which mechanisms are appropriate for minimizing a hull's size.

There are different ways to define maximal convexity and minimal size, and they yield generalized convex hulls with different geometric properties. Our goal is to define a hull with properties that are useful in low-dimensional geometric computing, particularly as applied to visualization, design, and manufacturing. In that context,

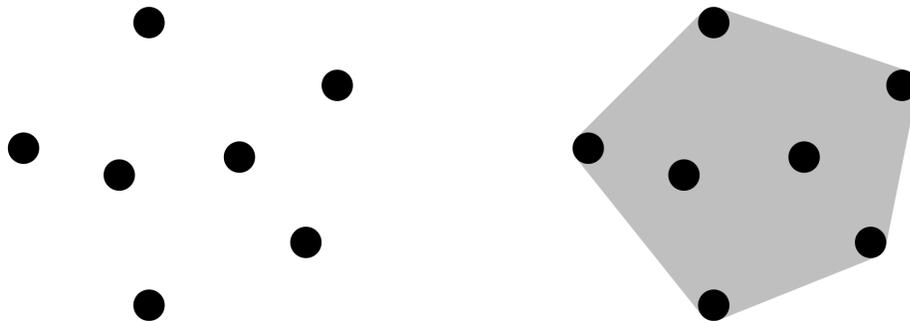


FIGURE 3.1. A set S consisting of isolated points is shown in black on the left. The convex hull of S is shown in gray on the right, with S superimposed.

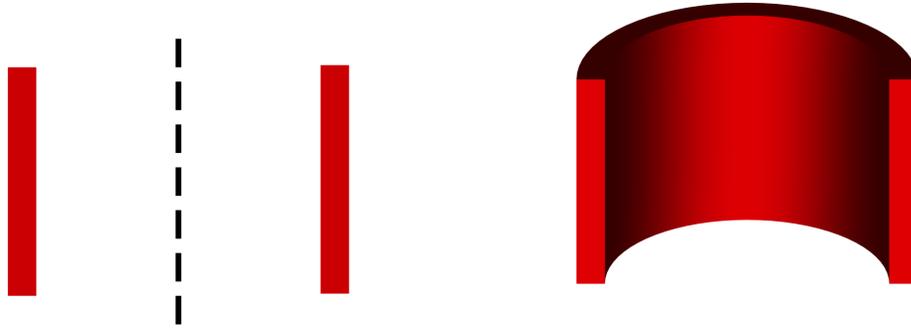


FIGURE 3.2. On the left we show a cross section of a hollow cylinder, which consists of two thin rectangles shown in red. If we rotate the rectangles about the dashed axis, they sweep the hollow cylinder, shown in a cutaway view on the right.

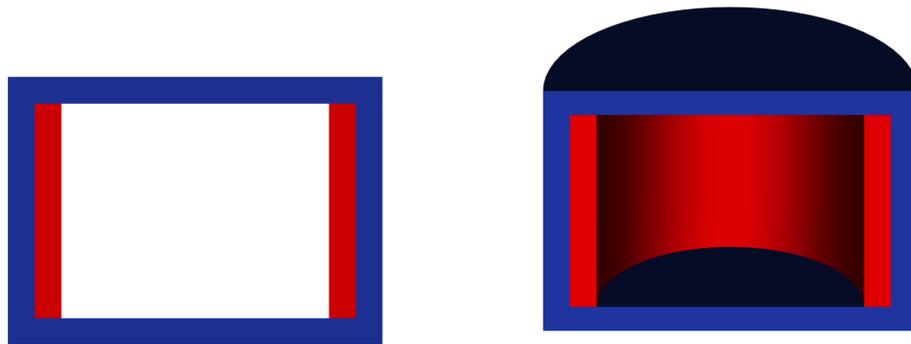


FIGURE 3.3. On the left, we show a cross section through the hollow cylinder's convex hull, with the hollow cylinder in red and the hull boundary in blue. We thicken the convex hull's boundary and offset it from the cylinder to aid visualization, but the curved portion of boundary of the hull actually lies on the the boundary of the cylinder. On the right, we show a cutaway view of the cylinder and its hull.

consider the following thought experiment: In one hand we hold a rigid, nonconvex object, with the boundary of the object's convex hull wrapped around it like a membrane. If we poke the membrane with a finger, how should it behave?

Suppose that the object is a hollow cylinder, like a section of pipe (Figure 3.2.) When wrapped in its convex hull, the hollow cylinder resembles a double-headed drum (Figure 3.3.) What happens when we press a finger against a drumhead, given that the drum remains maximally convex? One possibility is that the drumhead acts like the boundary of a fluid, surrounding the finger as it enters the object's convex hull

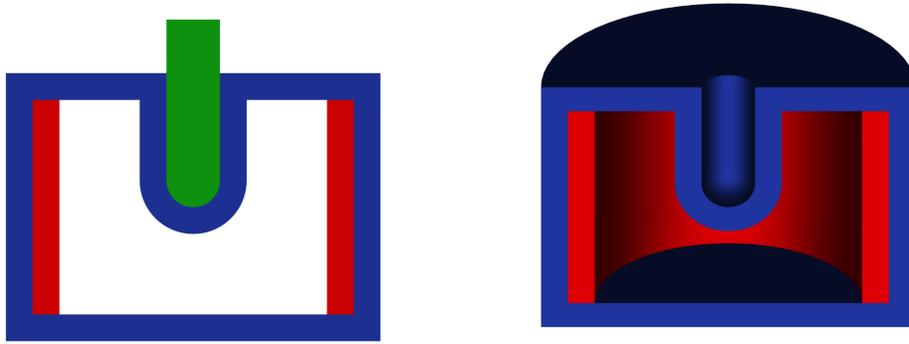


FIGURE 3.4. On the left, we show a cross section of a possible hull in blue whose boundary separates the hollow red cylinder from the interior of a green set corresponding to a finger pressing into the cylinder's convex hull. On the right, we show a cutaway of the three-dimensional hull. In both views, the hull surrounds the finger as if the hull boundary were fluid. In this respect, it behaves like a relative convex hull (Definition 8.11 [57, 58].)

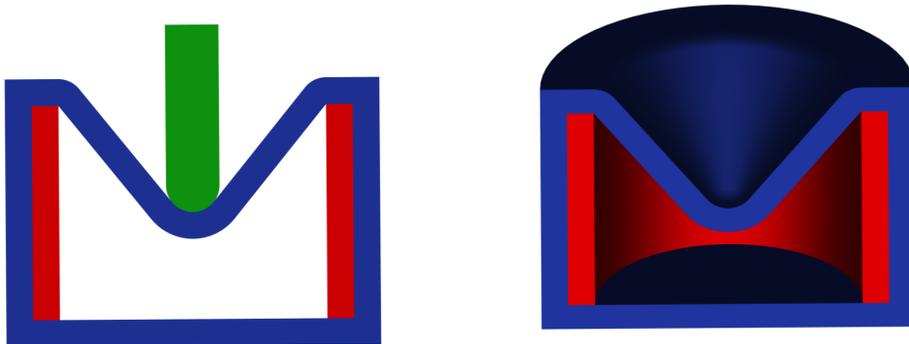


FIGURE 3.5. On the right, we show a cross section of a possible hull in blue whose boundary separates the red cylinder from a green set corresponding to a finger pressing into the cylinder's convex hull. On the right, we show a cutaway of the three-dimensional hull. In both views, the hull boundary behaves like a membrane in tension, stretched by the finger's pressure. In this respect, the hull behaves like a tight hull (Subsection 3.2.3, Chapter 13.)

(Figure 3.4.) The portion of the membrane that touches neither the finger nor the walls of the drum remains flat, and the finger causes a minimal change in the drum's volume. Alternatively, the membrane may behave as if it were in tension, in which case the finger deforms the drumhead into a cone (Figure 3.5.)

The relative convex hull, originally defined by Sklansky and Kibler [57], has a fluid-like boundary (Figure 3.4,) while the tight hull we introduce in this thesis has a boundary that behaves like a membrane in tension (Figure 3.5.) In two dimensions, the convex hull of R relative to G and the tight hull of R relative to G are not always identical (Figure 3.22,) but both sets have boundaries that behave like rubber bands. We formalize this behavior by defining a property called tightness in Chapter 13. Under our definition, both tight hulls and relative convex hulls are locally tight in two dimensions, while in higher dimensions tight hulls are tight and relative convex hulls generally are not. We argue that tightness can be a useful property, making tight hulls more suitable than relative convex hulls for some applications.

3.2. Comparison of hull definitions

3.2.1. Convex hull. In this section, we compare the definitions of a convex hull, a relative convex hull, and a tight hull. After defining each hull, we describe how to obtain it by successively filtering the set of all subsets of Euclidean space.

We adopt the following definition of the convex hull of a subset of Euclidean space [63]:

Definition 3.1. A set $S \subseteq \mathbb{R}^d$ is *convex* if and only if for every two points $p, q \in S$, the closed line segment \overline{pq} is contained in S (Figure 3.6.)

The *convex hull* of $S \subseteq \mathbb{R}^d$ is the intersection of all convex sets containing S (Figure 3.7.)

For every $S \subseteq \mathbb{R}^d$, a unique convex hull of S exists, and it is convex. The convex hull of S exists and is unique because \mathbb{R}^d is convex, so there is at least one convex set containing S , and there always exists a unique intersection of any nonempty collection of sets. The convex hull of S is convex because if it contains points p and q , every convex set containing S contains p and q . Because those sets are convex, each contains the segment \overline{pq} , so the convex hull of S contains \overline{pq} as well.

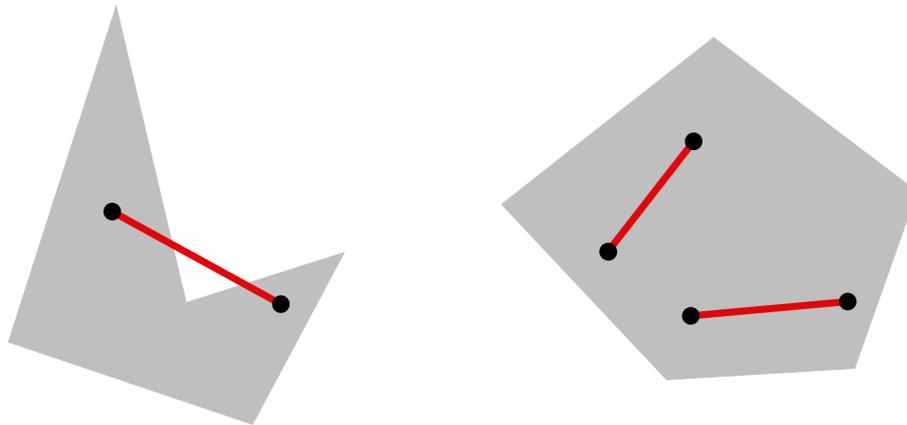


FIGURE 3.6. The gray polygon on the left is not convex because it contains the two black points but not the red segment connecting them. The polygon on the right is convex because it contains every segment that connects a pair of its points.

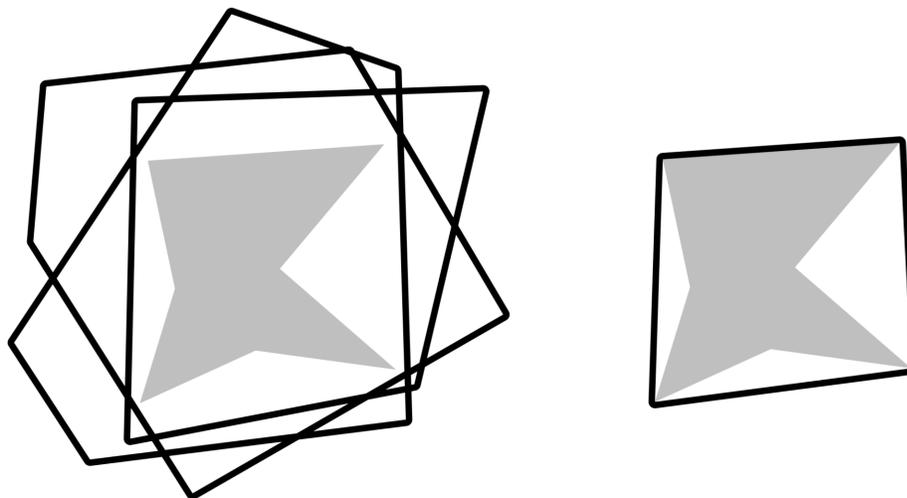


FIGURE 3.7. On the left we show a set S in gray and the boundaries of several of the convex sets containing S in black. On the right, we show the boundary of the convex hull of S in black. The convex hull of S is the intersection of the convex sets containing S .

To obtain the convex hull by filtering the set of subsets of Euclidean space, we successively remove those sets that lack one of the convex hull's properties. In the list of filtering steps below, we provide a name in parentheses for the property guaranteed by each step:

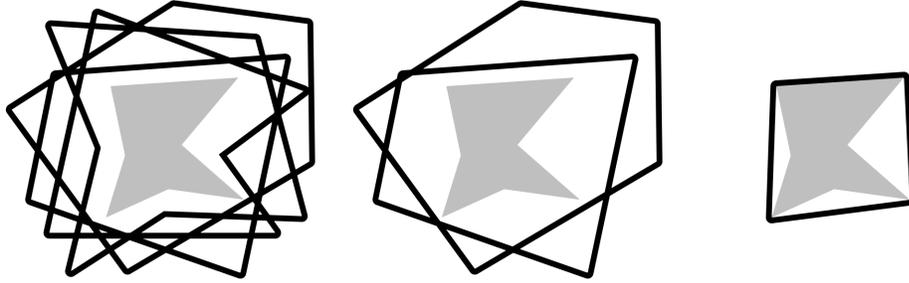


FIGURE 3.8. On the left, we show in gray the boundaries of several of the sets containing the set S shown in black. In the middle, we show the boundaries of convex sets containing S , and on the right we show the boundary of the smallest convex set containing S , which is the convex hull of S .

- (1) (Containment) Start with all subsets of Euclidean space that contain S (Figure 3.8, left.)
- (2) (Convexity) Extract the convex sets (Figure 3.8, middle.)
- (3) (Smallness) The convex hull of S is the smallest of the remaining sets when they are ordered by set inclusion (Figure 3.8, right.)

When the convex sets containing S are ordered by set inclusion, the convex hull is the smallest because, being equal to the intersection of the other sets, it is contained in all of them.

3.2.2. Relative convex hull. The definition of the convex hull of one set relative to another is similar to the definition of the Euclidean convex hull. We slightly modify the definitions provided by Sklansky and Kibler [57] and Sloboda and Zaitko [58] to facilitate our later analysis of symmetry with respect to set complement (Subsubsection 8.8.3.)

Definition 3.2. For $G \subseteq \mathbb{R}^d$, a set $S \subseteq \mathbb{R}^d$ disjoint from G is *convex relative to G* if and only if for every two points $p, q \in S$ such that \overline{pq} is disjoint from G , set S contains \overline{pq} (Figure 3.9.)

For $R \subseteq \mathbb{R}^d$ disjoint from $G \subseteq \mathbb{R}^d$, the *convex hull of R relative to G* is the intersection of all sets convex relative to G that contain R (Figure 3.10.)



FIGURE 3.9. The set S shown in gray on the left is not convex relative to the green set G because the red line segment connecting the two black points is not contained in S despite the fact that it is disjoint from G . The set shown in gray on the right, by contrast, is convex relative to G .

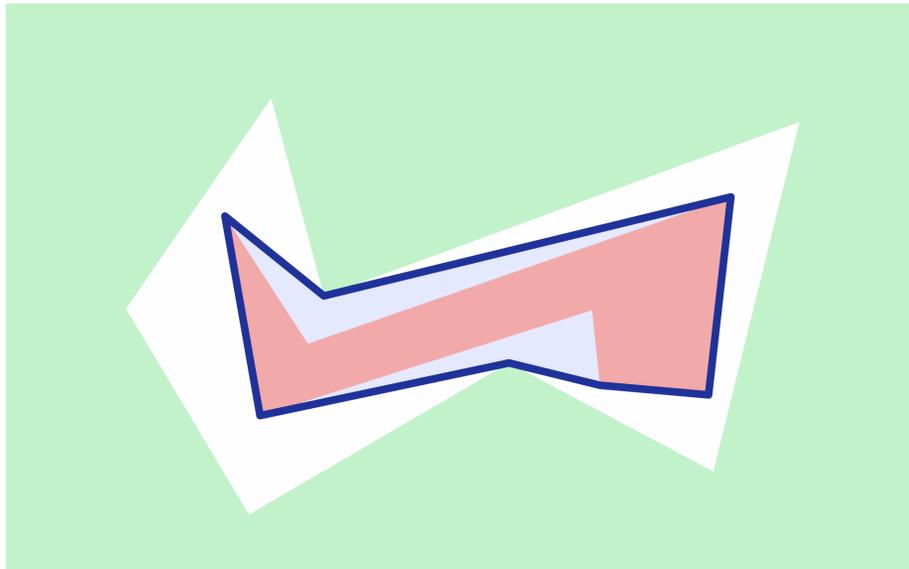


FIGURE 3.10. This image shows the boundary of the convex hull of R relative to G in dark blue, where R and G are shaded red and green, respectively. The hull contains R , and the material it adds to R is shaded light blue.

A unique convex hull of R relative to G exists, and it is convex relative to G . It exists and is unique because the complement of G contains R and is convex relative to G , so the intersection of all sets containing R that are convex relative to G exists and is unique. The convex hull of R relative to G is convex relative to G because if p and q are points in the convex hull of R relative to G , every set containing R that is

convex relative to G contains both p and q . If \overline{pq} is disjoint from G , all of those sets contain \overline{pq} , so the convex hull of R relative to G contains \overline{pq} as well.

We can obtain the convex hull of R relative to G by filtering the set of subsets of Euclidean space in a sequence of steps similar to those we used to obtain the convex hull:

- (1) (Containment) Start with all subsets of Euclidean space that contain R .
- (2) (Relative Convexity) Extract the sets that are convex relative to G .
- (3) (Smallness) The convex hull of R relative to G is the smallest of the remaining sets when they are ordered by set inclusion.

The definitions of the convex hull and relative convex hull are similar, the difference between them being that the convex hull is an intersection of convex sets, while the relative convex hull is an intersection of relatively convex sets. Convex sets and relatively convex sets are also similar, because the intersection of a collection of convex sets is convex, while the intersection of a collection of relatively convex sets is relatively convex. The same kinds of arguments that establish the convex hull's existence, uniqueness, and convexity then also establish the relative convex hull's existence, uniqueness, and relative convexity.

We can use the theory of convex structures [63], which we summarize in Section 8.7, to capture the similarities in behavior between the convex hull, the relative convex hull, and several similar sets. To do so, we start with a set X and a collection $\mathbb{C} \subseteq \wp(X)$ of subsets of X . We then define the sets in \mathbb{C} as convex if \mathbb{C} has certain properties, notably that the intersection of a subset of the elements in \mathbb{C} is always an element of \mathbb{C} . We define the convex hull of $Y \subseteq X$ as the intersection of all elements of \mathbb{C} that contain Y , and then we show that for every Y , a unique convex hull of Y exists and is convex.

When discussing subsets of Euclidean space, we use the terms “convex” and “convex hull” as we define them in Definition 3.1. When discussing convex structures, we use the terms “convex” and “convex hull” in the more inclusive senses of Definition



FIGURE 3.11. The round torus to the left is a tight embedding of a torus, while the ridged torus at the center and the warped torus to the right are not.

8.19, and use the term “Euclidean convex hull” to refer to the convex hull specified by Definition 3.1. We handle topological terms in a similar manner. A term like “closure” refers by default to the topological closure on \mathbb{R}^d defined by the Euclidean metric, but in the context of a discussion of closure structures (Subsection 8.6.1,) it refers to any closure operator.

3.2.3. Tightness and the tight hull. Tight sets do not define a convex structure, in part because tightness is formulated in terms of the boundary of a set rather than what it contains. We refer to a smooth surface as tight if it minimizes total absolute curvature for its topological type. (The total absolute curvature of a smooth surface is the integral over the surface of the absolute value of its Gaussian curvature, where the Gaussian curvature at a point on the surface is the product of the principal curvatures at that point.) Our definition of tightness is less restrictive than the use of the term in the theory of tight embeddings (Chapter 11; [13, 36, 69].) Any definition of tightness implicitly asks us to identify tight embeddings of different topological manifolds. For instance, we can establish that a round torus is a tight embedding of a topological torus (Figure 3.11,) while any smooth surface bounding a convex set is a tight embedding of a topological sphere (Figure 3.12.)

For a watertight triangle mesh in three dimensions, Gaussian curvature is concentrated at mesh vertices. Although the Gaussian curvature at a vertex is an impulse, the integral Gaussian curvature over any small neighborhood containing a vertex is



FIGURE 3.12. Both the sphere on the left and the rounded cube at the center are tight embeddings of a topological sphere. The shape on the right, which is the sweep of a sphere along a circular arc, is not tight.

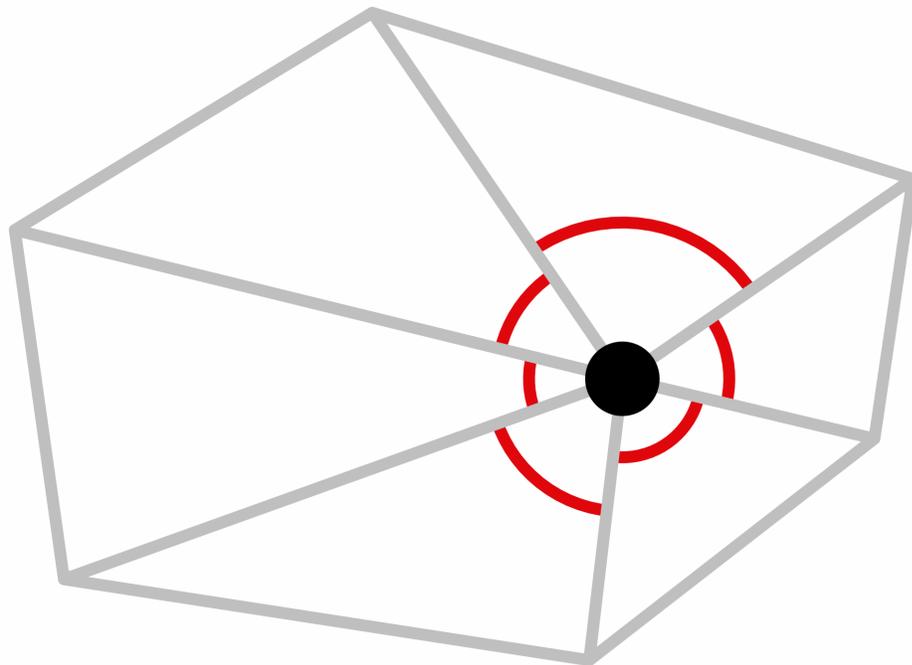


FIGURE 3.13. A vertex of a triangle mesh is shown in black with the triangles incident on it outlined in gray. The Gaussian curvature at the black vertex is 2π minus the sum of the angles subtended by the red arcs.

equal to the difference between 2π and the sum of the angles formed by consecutive edges incident on the vertex (Figure 3.13.)

Because the length of a small geodesic circle centered on the vertex is proportional to the sum of its incident angles, the Gaussian curvature at a point on a surface in three dimensions measures the difference between the length of a circle in the plane

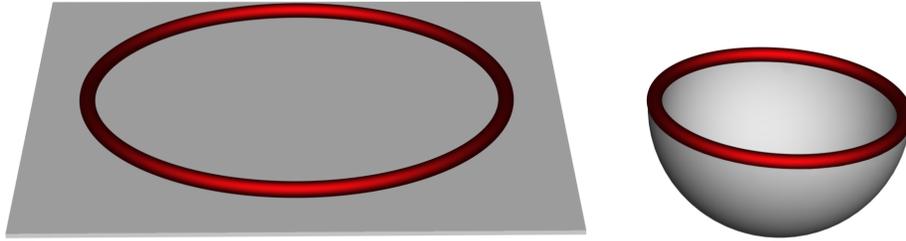


FIGURE 3.14. The surface on the left is planar, while the surface on the right is half of a unit sphere. The curves shown in red are geodesic circles of radius $\frac{\pi}{2}$. Because the surface on the right is curved, the geodesic circle on it has length 2π , while the circle on the left has length π^2 . The fact that $\pi^2 > 2\pi$ reflects the fact that spheres have positive Gaussian curvature. The value of the Gaussian curvature is equal to one at every point on the hemisphere, because every sectional curvature at a point on the hemisphere is one. If we let $r \in \mathbb{R}^>$ denote the radius of both geodesic circles, we can compute that value by taking the limit as r goes to zero of the difference between the lengths of the geodesic circles multiplied by $\frac{3}{\pi r^3}$ [4].

and the length of a geodesic circle on the surface centered at the point when we the two circles' radii are small and equal (Figure 3.14.) At a saddle, the length of the geodesic circle is greater than the length of the flat circle, so the Gaussian curvature is negative, while at a vertex with a convex or concave neighborhood the length of the geodesic circle is smaller, so the Gaussian curvature is positive. (We distinguish between vertices whose neighborhoods are convex or concave, and convex or concave points of a set, which locally maximize distance from some plane over the set (Section 12.2.) A convex vertex may have negative Gaussian curvature, while a the Gaussian curvature at a vertex with a convex neighborhood is nonnegative.)

To compute the total absolute curvature for a mesh, we obtain polyhedral absolute Gaussian curvature values as described in Subsection 11.2.3. The total absolute curvature of a mesh is then the sum of these values at all mesh vertices. Under this definition, every convex polyhedron is tight, and the total absolute curvature of a convex polyhedron under the polyhedral definition is equal to the total absolute curvature of a sphere under the smooth definition (Figure 3.15.)

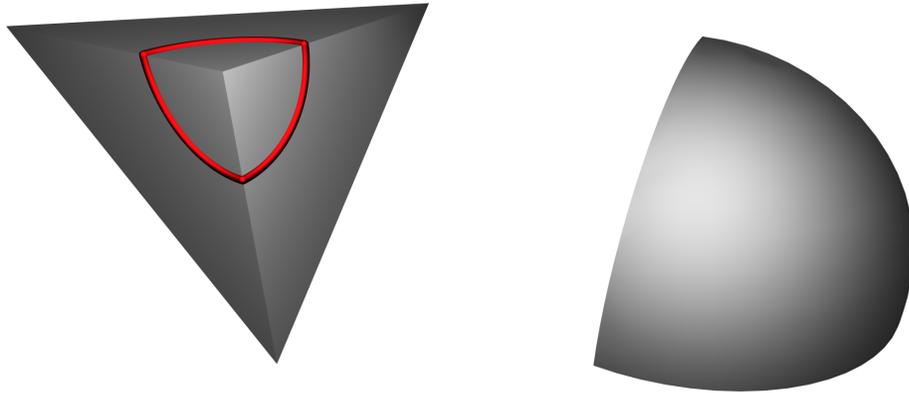


FIGURE 3.15. On the left we show a corner of a cube. No matter how we triangulate the cube's boundary, the sum of the angles incident on the vertex at one of its corners is equal to $\frac{3\pi}{2}$. Under our definition of polyhedral Gaussian curvature, the Gaussian curvature at a corner is then $2\pi - \frac{3\pi}{2} = \frac{\pi}{2}$. Because there are eight corners, the cube's total absolute curvature is 4π . On the right we show an octant of a sphere, which like the corner of a cube contributes $\frac{\pi}{2}$ to the sphere's total absolute curvature. Like a cube, a sphere's total absolute curvature is 4π . This is true regardless of the value r of its radius, because the sphere's surface area is $4\pi r^2$, while its Gaussian curvature is $\frac{1}{r^2}$.

In three dimensions, we can unify smooth and polyhedral tightness using the two-piece property. A subset of d -dimensional Euclidean space has the two-piece property if and only if its intersection with a halfspace has no more than one connected component (Definition 11.17;) equivalently, the result of subtracting the boundary of any halfspace from the set has no more than two connected components. A sphere, a ball, a convex polytope, a round torus, and various polygonal torii all have the two-piece property, and all of them are tight (Figure 3.16.)

We do not, however, define tightness in terms of the two-piece property. Instead, we define tightness in terms of a measure of normal variation called slack (Definition 10.14) that extends the definition of total absolute curvature to arbitrary subsets of Euclidean space. We use slack in part because in two and higher dimensions, there exist sets R and G such that no set that contains R and excludes G has the two-piece property (Figure 3.17.) This prevents us from directly applying the two-piece property to our definition of a tight hull. Moreover, in higher dimensions slack

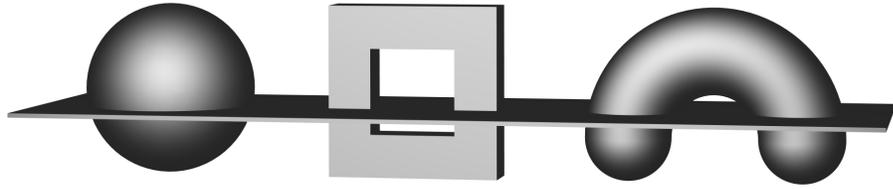


FIGURE 3.16. A sphere and a polygonal torus have the two-piece property, as illustrated on the left and in the center. The set on the right is cut into three connected components by a plane, so it is not tight.

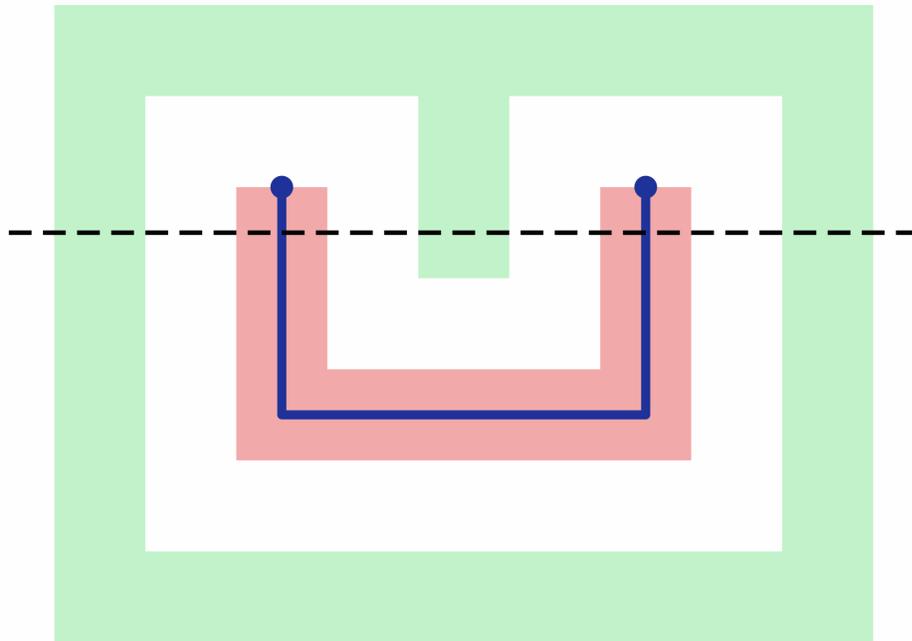


FIGURE 3.17. In the two-dimensional example above, no set that contains R and excludes G has the two-piece property. A set that contains R must contain the two points on R indicated by blue dots, but any path connecting those two points must cross the dashed black line, implying the points are in two different connected components when the hull is cut by the line. Because a portion of R also lies below the line, the cut yields at least three pieces.

minimization is a stronger condition than the two piece property: sets with the two-piece property do not necessarily minimize slack, but if an embedding of a manifold with the two-piece property exists, a slack-minimizing embedding of the manifold has the two-piece property. We therefore broadly characterize tight sets as slack

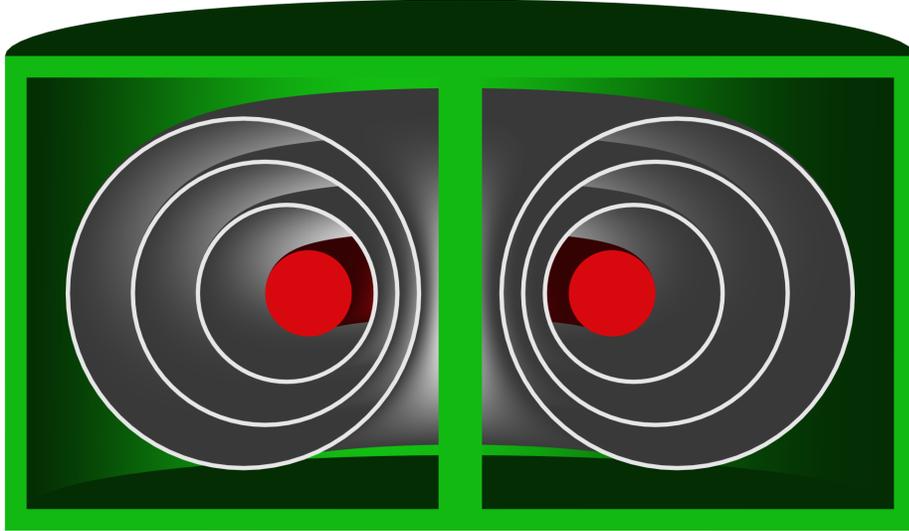


FIGURE 3.18. In this figure, R is the red torus and the complement of G is a punctured cylinder. All of the gray torus boundaries shown separate R and G with minimal slack. Because the radii of the circular torus cross sections can vary continuously, there are uncountably many torii of this kind.

minimizers. This characterization remains meaningful regardless of spatial dimension or the constraints on the slack minimization.

Under modest assumptions, there is always a slack minimizer that contains R and excludes G , and in general there are uncountably many (Figure 3.18.) Taking these constrained slack minimizers as candidate tight hulls, we define the tight hull as a candidate tight hull that additionally minimizes unsupported slack. To define unsupported slack, we first define the convexity of an outward-pointing unit normal $v \in \mathbb{S}^{d-1}$ to $S \subseteq \mathbb{R}^d$ at $p \in \mathbb{R}^d$ as the fraction of two-dimensional sections of S containing the ray with tail p and direction v that are locally convex at p . Convexity is then a value in $[0, 1]$. A convexity normal to S at p is a normal to S at p scaled by its convexity.

The support of a convexity normal v to S at p from $T \subseteq S$ is the maximum value over all dot products of v with the set containing the zero vector and the convexity normals to T at p . Support, like convexity, is a value in $[0, 1]$ (Figure 3.19.) The support of a subset of the boundary of S from T is the integral of the support of its

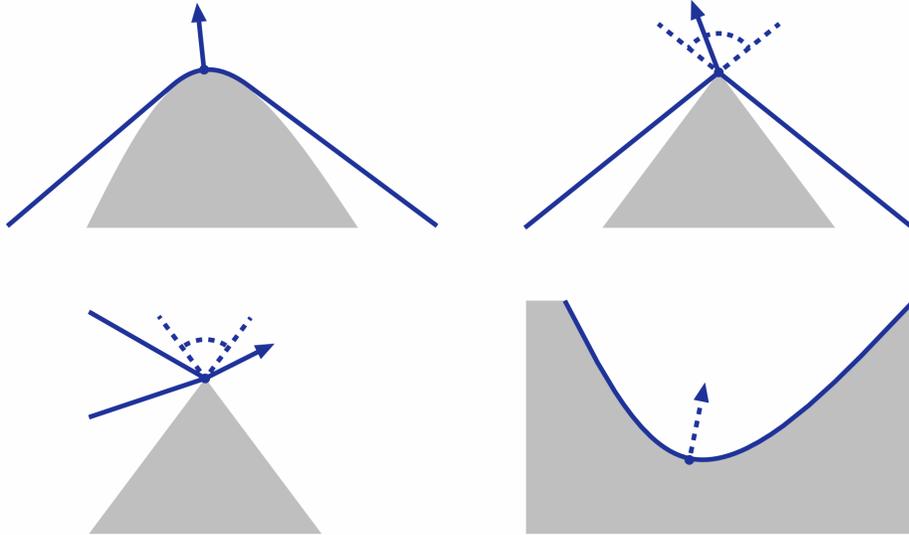


FIGURE 3.19. In these figures, the hull boundary and normals to the hull are shown in blue, while the constraints are shaded gray. At the top left, the normal to the hull at the marked point is equal to the normal to the constraints, so its support is one. At the top right, the normal to the hull is an element of the fan of normals at the corner, so its support is also one. At the lower left, the hull normal drawn is outside of the fan of normals at the corner, but it has a positive dot product with a normal to the constraints at the corner. As a result, its support is between zero and one. At the lower right, the convexity normal at the marked point has zero length, so the hull normal to the candidate hull is unsupported.

outward-pointing normals, which is a value between zero and the subset's slack. The support of a candidate tight hull of R relative to G is the sum of the hull's support from R and the support of the hull's complement from G . A candidate tight hull's unsupported slack is the difference between its slack and its support (Figure 3.20.) If the support of a subset of a candidate hull's boundary is zero, we refer to the subset as unsupported. We provide a more formal definition of unsupported slack in Section 12.2.

We now define the tight hull, letting X^c denote the complement of a set X :

Definition 3.3. For $R, G \subseteq \mathbb{R}^d$, suppose $R \cap G = \emptyset$. Then $C \subseteq \mathbb{R}^d$ is a *candidate tight hull* of R relative to G if and only if

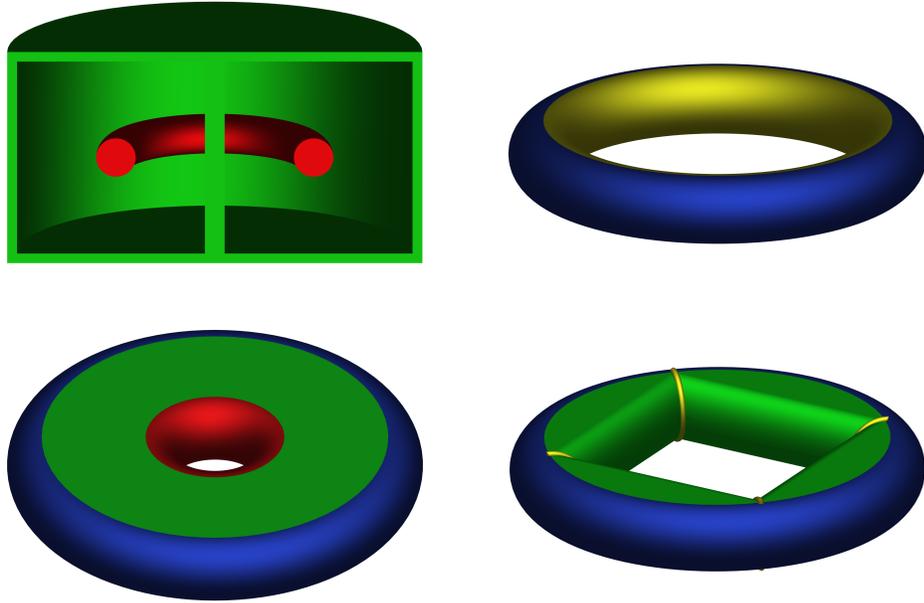


FIGURE 3.20. On the top left, we show sets R and G identical to the R and G in Figure 3.18. The tight hull of R relative to G is equal to R . We show the tight hull on the top right, with the boundary colored to indicate the amount of support at each point. The blue part of the boundary on the outer wall of the torus is fully supported, while the yellow part is only partially supported, because some sections at the yellow points are convex and others are concave. On the lower left, we show a candidate hull that minimizes slack but not unsupported slack. The red portion of the hull boundary is unsupported. The green portion is also unsupported, but it is developable, so its slack is zero and it does not contribute unsupported slack. We show another candidate hull on the bottom right. Part of the hull's slack is supported on the four yellow arcs, but the support from those arcs is less than the support from the inner wall of the torus in the tight hull on the top right, so this candidate is not a tight hull.

- (1) $R \subseteq C$
- (2) $C \subseteq G^c$
- (3) There is no $S \subseteq \mathbb{R}^d$ such that $R \subseteq S$, $S \subseteq G^c$, and the slack of S is less than the slack of C .

A candidate tight hull T of R relative to G is a *tight hull* of R relative to G if and only if there is no candidate tight hull S of R relative to G such that the unsupported slack of S is less than the unsupported slack of T .

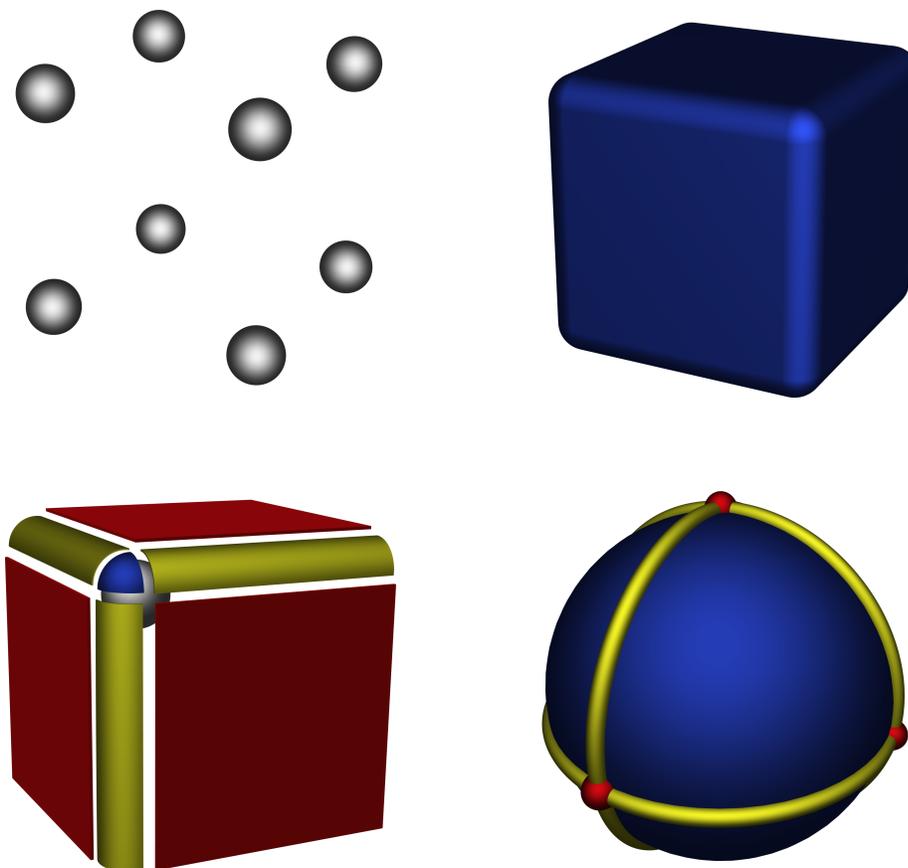


FIGURE 3.21. At the top left, we show eight small spheres, each centered on the corner of a unit cube. At the top right, we show the convex hull of the spheres in blue. At the bottom left, we show an exploded view of one of the hull's corners, with the hull decomposed into a blue spherical cap, yellow cylindrical edges, and planar red faces. On the bottom right, we show use the same colors to indicate which of the normals to the hull can be attributed to its spherical, cylindrical, and planar components. The blue triangles are due to the hull's spherical caps, which are supported. The union of the yellow arcs and red points is a lower-dimensional set of unsupported normals due to the hull's edges and faces. Because this set has zero measure, the hull's normals are fully supported.

Defined in this way, the tight hull generalizes the convex hull: the convex hull has minimal slack because it is convex, and all of its slack is supported (Figure 3.21; Theorem 12.9.) The procedure for constructing a tight hull is:

- (1) (Separation) Start with all sets that contain R and exclude G .
- (2) (Tightness) Extract the sets that minimize slack.

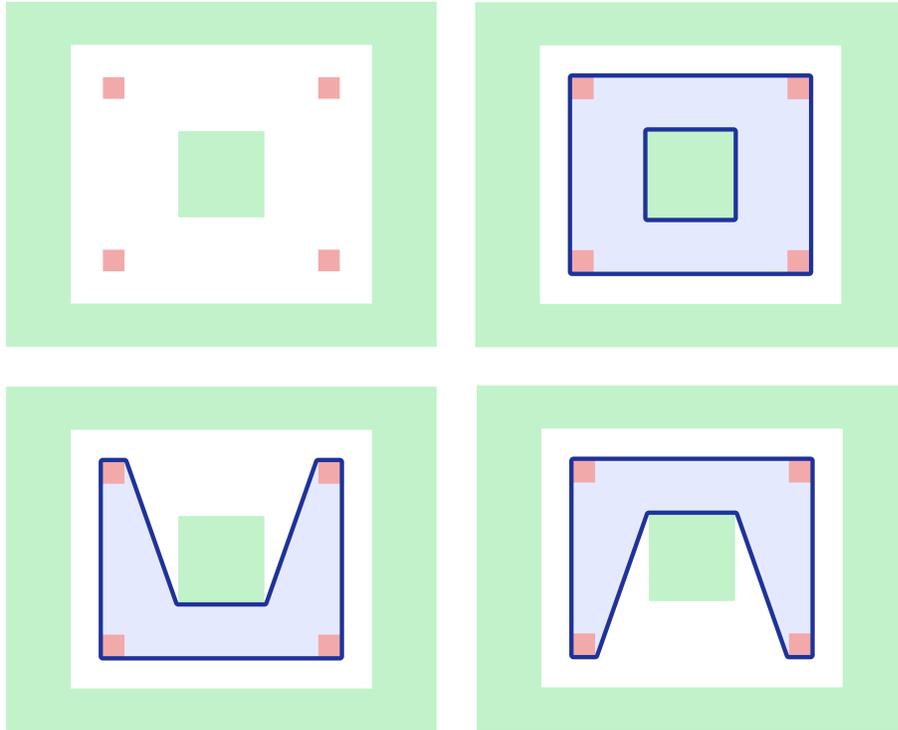


FIGURE 3.22. We show the sets R and G in the upper left. The set R consists of four red squares, while the complement of G is a rectangle with a square hole. We show the convex hull of R relative to G in the upper right; note that the hull does not minimize slack, because the sum of the angles around its two bounding loops is 4π . The two tight hulls of R relative to G appear on the lower right and left. Both hulls have equal slack (roughly 3.5π), and their slack is fully supported. Although both are tight hulls, their interiors differ.

(3) (Smallness) Extract the sets that minimize unsupported slack.

A tight hull always exists, but its interior may not be unique (Figure 3.22.) We tentatively attribute nonunique hull interiors to singular configurations of R and G , and we conjecture that we can obtain a unique hull by perturbing either.

3.2.4. Comparison of tight and relative convex hulls. When we compare the definitions of the relative convex hull and tight hull, we see that they both contain R and exclude G ; in the relative convex hull, the exclusion of G is due to the relative convexity requirement. Both hulls also maximize convexity, but in different ways. While a relative convex hull contains as many of the line segments connecting pairs

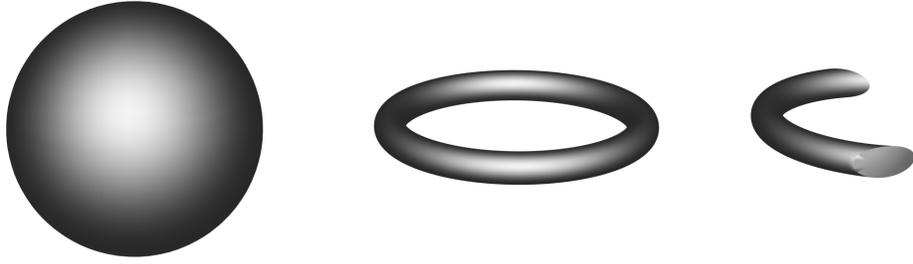


FIGURE 3.23. The sphere on the left and torus in the center are tight, but their intersection on the right is not.

of its points as possible, a tight hull minimizes slack. Because the definitions maximize convexity differently, they minimize size differently. Ordering the relatively convex sets by inclusion is appropriate because the intersection of relatively convex sets is relatively convex.

The intersection of tight sets, however, is not always tight (Figure 3.23.) Consequently, we cannot define a convex structure consisting of tight sets and obtain a tight hull from it. The slack minimizers out of all sets that contain R and exclude G can often be deformed into each other, and such deformations redistribute slack along the deformed set's boundary. Minimizing a hull's unsupported slack concentrates slack on the subset of the hull boundary that touches R and G . We therefore conjecture that the boundary of the tight hull is developable where it does not touch R and G , which is a property shared by the convex hull and relative convex hull.

If true, this implies that the tight hull of one three-dimensional polyhedron relative to another is polyhedral, so we can exactly represent it with a triangle mesh. More generally, our hope is that if we model R and G with a concise, precise representation, we can also represent the tight hull of R relative to G concisely and precisely.

3.3. Hulls of revolution

We now explain the principal geometric differences between tight hulls and relative convex hulls, focusing on those differences that emerge when we move from two dimensions to three dimensions. We begin with planar sets R and G that both have

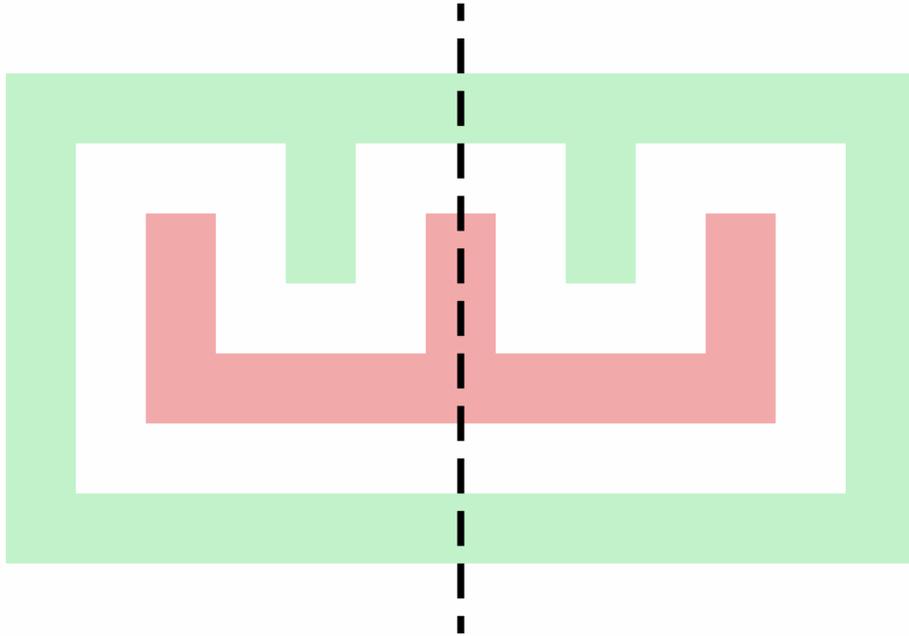


FIGURE 3.24. The sets R and G are shown in red and green, respectively. Both R and G^c are simple polygons symmetric about the dashed black line l .

mirror symmetry about line l (Figure 3.24.) Both R and the complement of G are simple polygons.

The tight hull of R relative to G and the convex hull of R relative to G are identical in this case (Figure 3.25,) and we show in Section 12.3 that the tight hull and relative convex hull are always identical when R is contained in G^c and both are homeomorphic to a disk. This result does not generalize to higher dimensions, however. Consider rotating R and G about l so that they sweep out solids R' and G' (Figure 3.26.) The set R' swept by R is the union of a hollow cylinder and a narrow cylinder with a wide, thin cylindrical base (Figure 3.27.) Similarly, G' is the union of a hollow cylinder with the unbounded complement of a cylinder (Figure 3.28.)

Figure 3.29 shows the convex hull of R' relative to G' . The outer walls of the hollow cylindrical portion of G' appear to press into the hull, while the hull boundary appears stretched over the narrow cylindrical part of R' near l . The hull's unsupported boundary is locally convex, which is a property of all relative convex hulls. As we

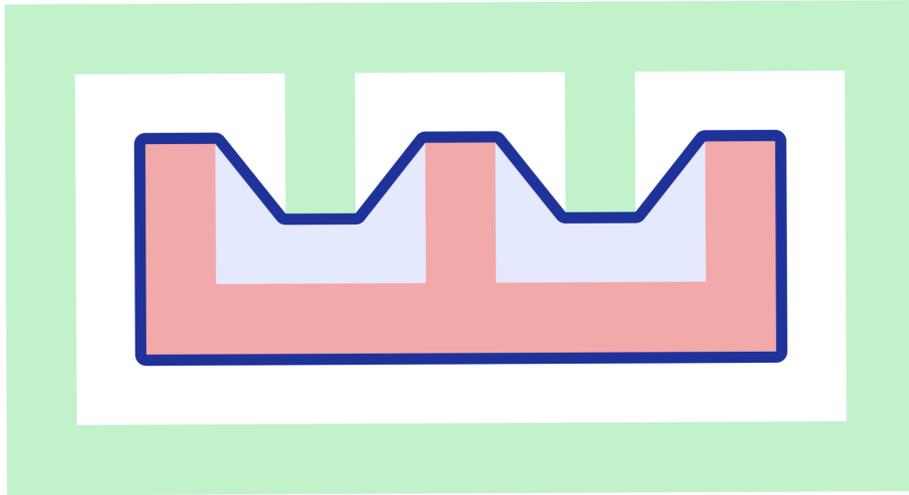


FIGURE 3.25. For the red set R and green set G shown, the convex hull of R relative to G and the tight hull of R relative to G are identical. The hull boundary is shown in dark blue, and the material the hull adds to R is shaded light blue.

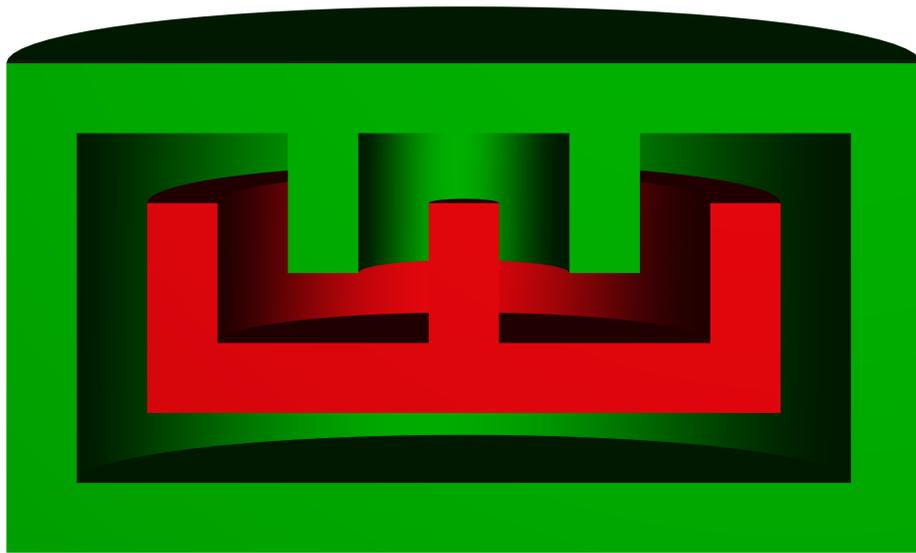


FIGURE 3.26. The red set R' and green set G' produced by rotating the planar sets R and G in Figure 3.24 about their line of symmetry l .

explain in Section 12.3, if there is an unsupported point on the boundary of a relative convex hull that is not locally convex, then there is a line segment connecting points of the hull that is disjoint from G but not contained in R . The convex hull of R relative to G is then not convex relative to G , which is a contradiction.

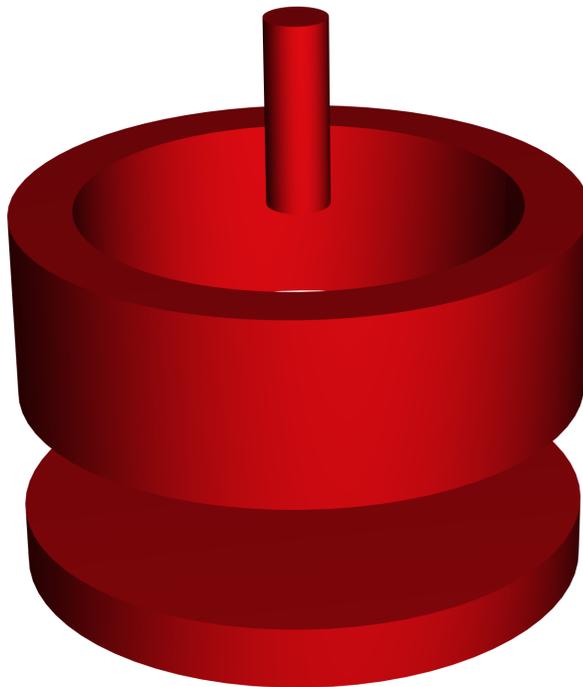


FIGURE 3.27. The set R' produced by rotating the set R shown in Figure 3.24 consists of a hollow cylinder on a wide cylindrical base with a narrow cylinder along their common axis. We show the three cylindrical pieces separately in this exploded view.

Figure 3.30 shows the tight hull of R' relative to G' . Qualitatively, the hull boundary is stretched over both R' and G' , and the unsupported boundary combines locally convex and locally concave regions. Because it contains unsupported concave points, the tight hull is not convex relative to G' .

When we compare cross sections of the three-dimensional tight hull and relative convex hull, we see that the section of the tight hull of revolution is equal to both the planar tight hull and the planar relative convex hull, while the section of the relative convex hull of revolution does not match either (Figure 3.31.) The two angled segments in the relative convex hull section sweep out the conical portion of the hull's boundary when rotated about l , while the segments orthogonal to l sweep a flat, annular patch. The angle that a segment from a two-dimensional hull makes with l determines whether the surface it sweeps is flat or conical, and whether the conical surface is convex or concave. Because the unsupported boundary from the

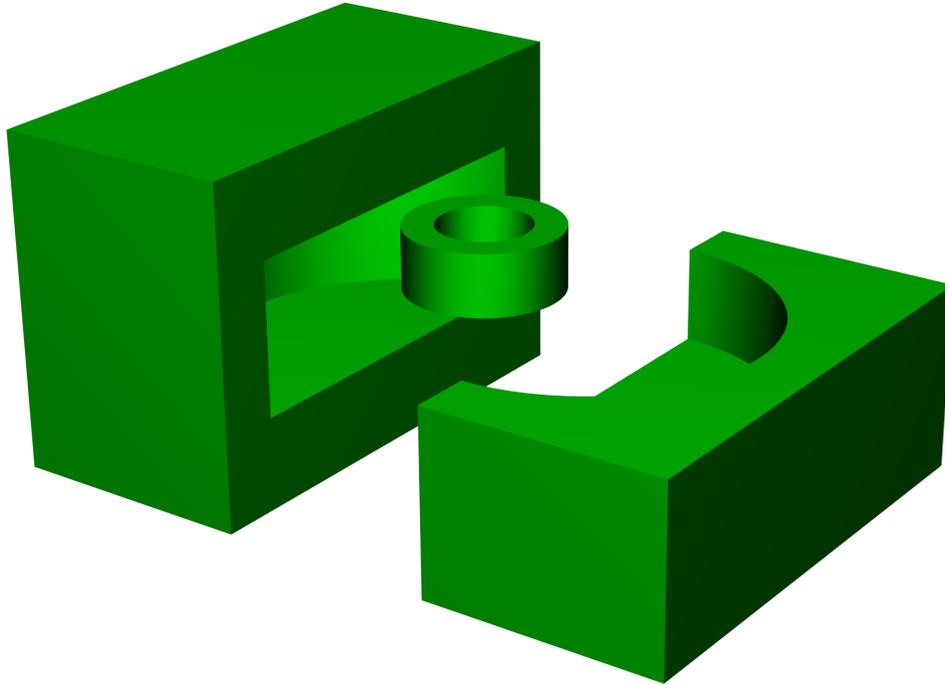


FIGURE 3.28. The set G' produced by rotating the set G shown in Figure 3.24 consists of the union of a hollow cylinder with the unbounded complement of a larger cylinder. We show an exploded cutaway of G' to make the hollow cylinder visible.

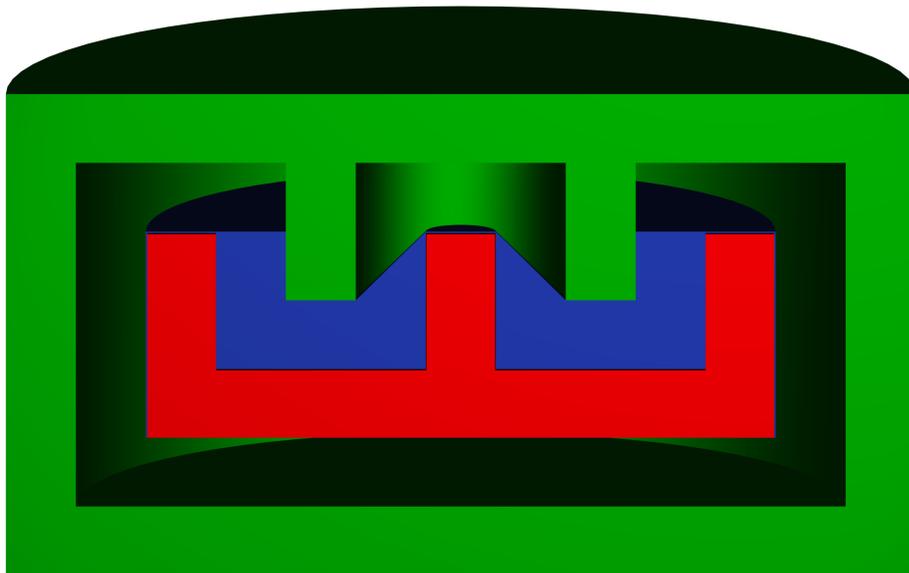


FIGURE 3.29. This figure shows a cutaway of the convex hull of R' relative to G' , where R' and G' are the red and green sets shown in Figure 3.26. The relative convex hull contains R' , and the material the hull adds to R' is shown in blue.

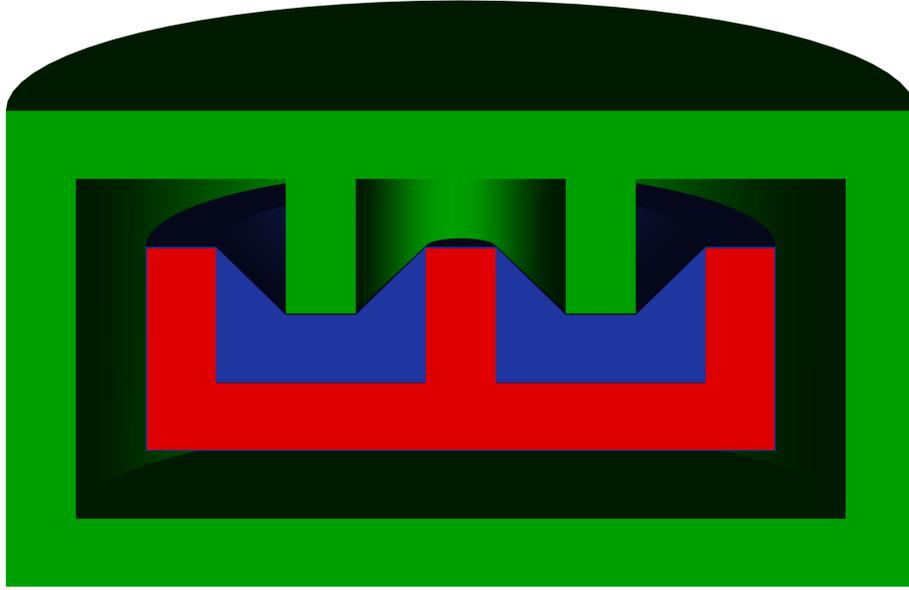


FIGURE 3.30. This image shows a cutaway of the tight hull of R' relative to G' in blue, where R' and G' are the red and green sets shown in Figure 3.26. The tight hull contains R' , and a cross section of R' is visible at the cutting plane.

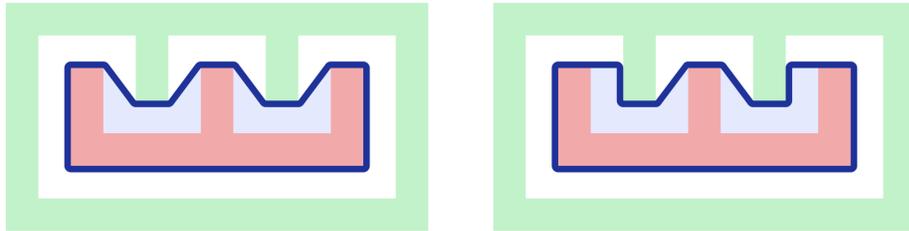


FIGURE 3.31. On the left we shown a cross section of the the three-dimensional tight hull depicted in Figure 3.30. The section is identical to both the tight hull and convex hull of the planar sets R relative to G shown in Figure 3.25, so that three-dimensional tight hull is the equal to sweep of the two-dimensional tight hull as it rotates about its axis of symmetry. On the right we show a section of the three-dimensional relative convex hull depicted in Figure 3.29. The cross section of the relative convex hull of the sets swept by rotating R and G is not itself a tight or convex hull of R relative to G .

relative convex hull is locally convex, the unsupported segments in one of its cross sections can only form a limited range of angles with l . There is no such restriction on the unsupported segments from a section of the tight hull.

Because our three-dimensional tight hull and relative convex hull are rotationally symmetric, their slack increases monotonically with the slack of their cross sections. The hull cross sections are polygons, and the total absolute curvature of a polygon is the sum of the absolute values of the change in normal direction at each vertex. The sum for the tight hull section from this example is 4π , while the sum for the relative cross hull section is 5π . We infer that the three-dimension relative convex hull has greater slack than the three-dimensional tight hull, so the relative convex hull does not minimize slack.

To establish that rotating our two-dimensional tight hull sweeps our three-dimensional tight hull, we show that if R and G are rotationally symmetric about the same axis l , the tight hull of R relative to G is also rotationally symmetric about l . Its slack and unsupported slack are then monotonically increasing functions of the slack and unsupported slack of its sections. A section that first minimizes slack and then minimizes unsupported slack over slack minimizers yields a solid that first minimizes slack and then minimizes unsupported slack over slack minimizers. We conclude that each section through the solid's axis of rotational symmetry must be a tight hull (Chapter 13.)

3.4. Symmetry with respect to set complement

Using the same swept solids R' and G' , we can construct the tight hull and convex hull of G' relative to R' (Figure 3.32.) These hulls contain G' and exclude R' , and so are unbounded. To directly compare them to the tight hull and convex hull of R' relative to G' , we take their set complements (Figure 3.33.) The images of the tight hull of R' relative to G' and the complement of the tight hull of G' relative to R' are indistinguishable, while the corresponding relative convex hull images are substantially different.

The unsupported boundary of a relative convex hull is locally convex, so the unsupported boundary of the complement of a relative convex hull is locally concave.

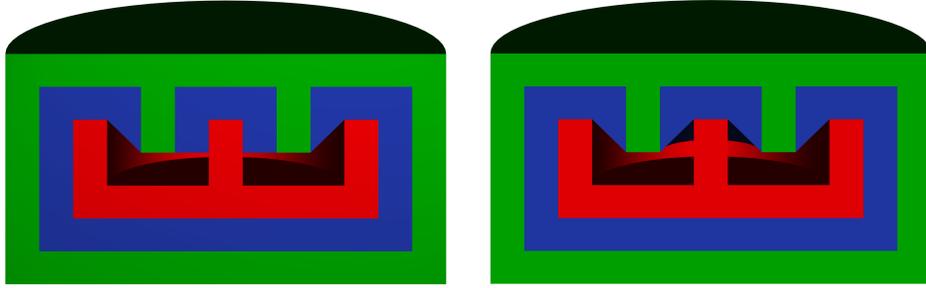


FIGURE 3.32. On the left, we show a cutaway of the convex hull of G' relative to R' . As before, the hull's unsupported boundary is locally convex. On the right, we show the corresponding tight hull. The unsupported boundary of the tight hull has both locally convex and locally concave components.

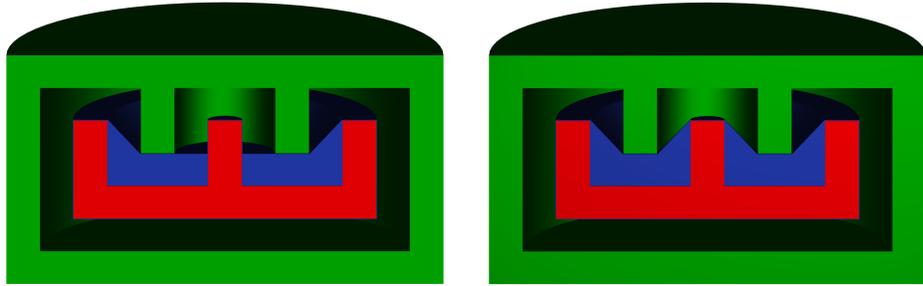


FIGURE 3.33. The set in blue on the left is the complement of the convex hull of G' relative R' . It appears substantially different from the convex hull of R' relative to G' shown in Figure 3.29. In particular, its unsupported boundary has nonpositive curvature, while the curvature of the unsupported boundary of the hull in Figure 3.29 is nonnegative. The figure on the right is the complement of the tight hull of G' relative to R' . It is indistinguishable from the tight hull of R' relative to G' shown in Figure 3.30.

The two are identical only if all of the hull's unsupported boundary is flat, which is not generally the case. We therefore characterize the relative convex hull as asymmetric with respect to set complement: the interior of the convex hull of R relative to G is not in general the exterior of the convex hull of G relative to R .

The tight hull, by contrast, is symmetric with respect to set complement (Chapter 13.) This is because the definition of a tight hull is formulated in terms of its boundary, and the boundary of a set is identical to the boundary of its complement. While taking the complement of the tight hull of R relative to G maps the hull's interior to the

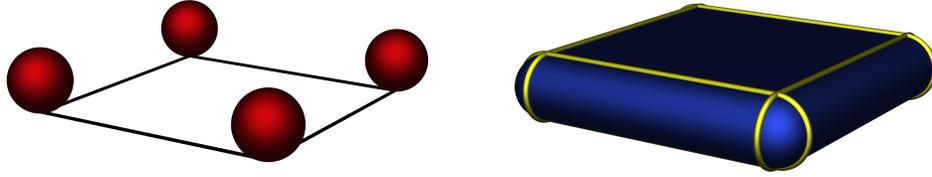


FIGURE 3.34. On the left we show four red balls of equal radius resting on the corners of a black square. The convex hull of the spheres, shown in blue on the right, consists of two squares, four cylindrical patches, and four spherical caps. We highlight the patch boundaries with yellow curves.

interior of the complement of the tight hull of G relative to R , our definition does not specify which subset of its boundary a tight hull includes, making our theorems concerning tight hulls independent of a tight hull's intersection with its boundary.

3.5. The tablecloth problem

We now present a more complex three-dimensional example that is not rotationally symmetric. Suppose R consists of four red balls of radius r resting at the corners of a square lying in a horizontal plane. The convex hull of R consists of two squares, four cylindrical patches, and four spherical caps (Figure 3.34.) From the top, the hull boundary resembles a tablecloth draped over the balls, but without any folds near the spherical caps. As in our drum example in Subsection 3.1, we can envision pressing a ball into the tablecloth at the center of the square patch on the top of the balls' convex hull.

Suppose that G is a green ball of radius r with its center at a distance of $2r$ above the point at the center of the square on the flat surface (Figure 3.35.) Figure 3.36 shows the convex hull of R relative to G , while Figure 3.37 shows the tight hull of R relative to G . As in our previous examples, the relative convex hull behaves like a fluid, surrounding the red ball. The red ball introduces a circular hole of radius r into the square patch on the top of the convex hull of G and fills the hole with a hemisphere.

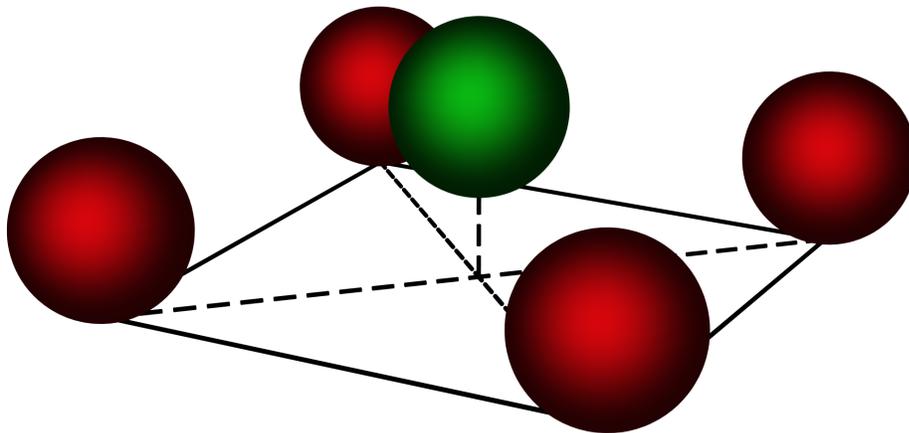


FIGURE 3.35. The set R consists of the four red balls of radius r arranged at the corners of a square. The set G is the green ball with its center at a distance of $2r$ above the square's center.

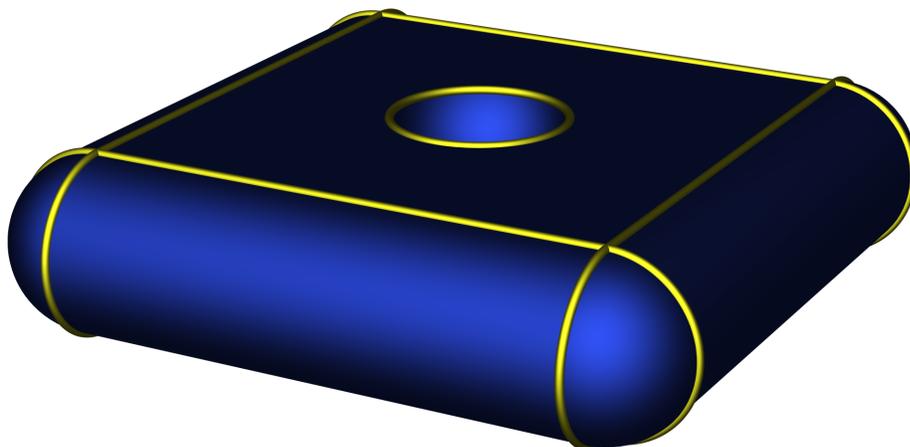


FIGURE 3.36. We show the convex hull of R relative to G in blue. The hull boundary consists of round cylinders, spherical caps, and a square with a circular hole. We highlight the patch boundaries in yellow.

The tight hull replaces the square on top of the convex hull of G with four triangles, four strips from elliptical cylinders, four conical patches, and a spherical cap. The cones, elliptical cylinders, and triangles meet at four singular vertices, while the elliptic patches meet along four singular arcs. The hull boundary is not smooth, despite the fact that the boundaries of R and G are smooth.

We can think of each elliptic cylinder in the tight hull boundary as being produced by the sweep of a circle along a line that forms an angle of $\frac{\pi}{2}$ with the plane containing

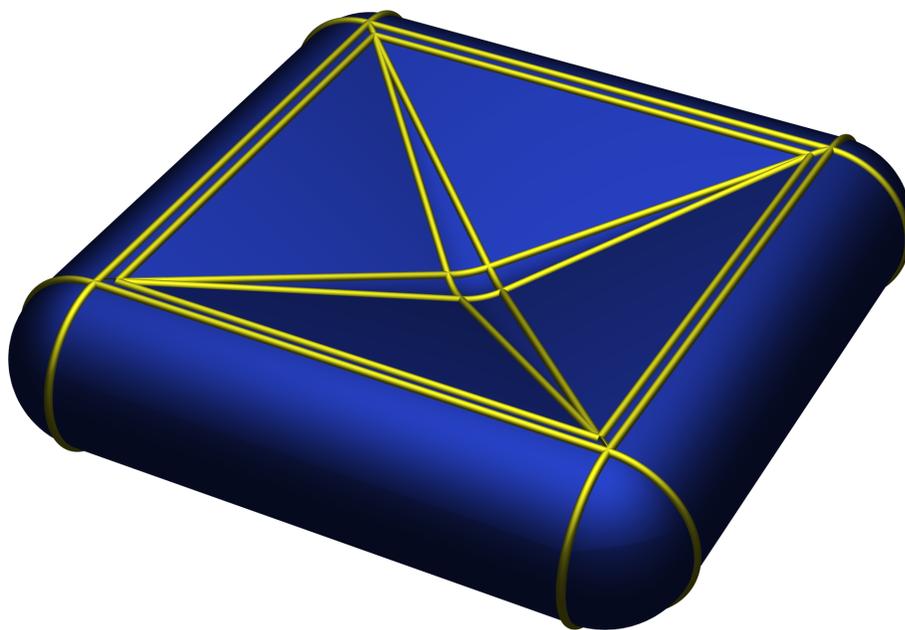


FIGURE 3.37. We show the tight hull of R relative to G in blue. The hull boundary consists of round cylinders, elliptical cylinders, spherical caps, and triangles. We highlight patch boundaries in yellow.

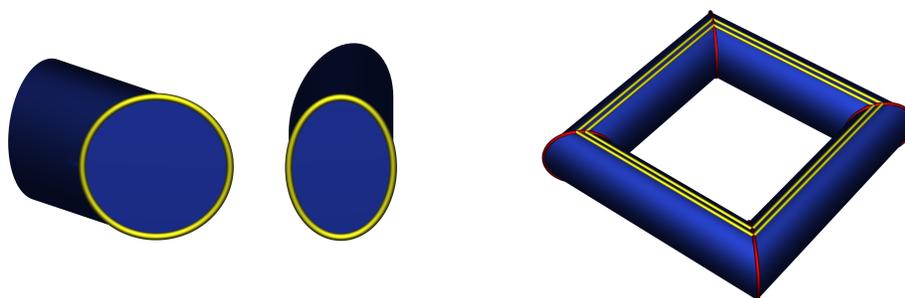


FIGURE 3.38. On the left, we illustrate how an elliptic cylinder is swept by a circle as it moves along a line that forms an angle of $\frac{\pi}{2}$ with the plane containing the circle. The cross section orthogonal to the sweep direction is elliptical. On the right, we show the elliptic cylinders swept by great circles from the red balls aligned with the diagonals of the square the balls form. We show the circles in red, and we outline the portion of the union of the cylinders that contributes to the tight hull boundary in yellow.

the circle. (Figure 3.38.) To explain the elliptic cylinders' presence, consider the similar, but incorrect, tight hull shown in Figure 3.39. We can construct this hull by first taking the union of the convex hulls of pairs of nondiagonal red balls, forming a

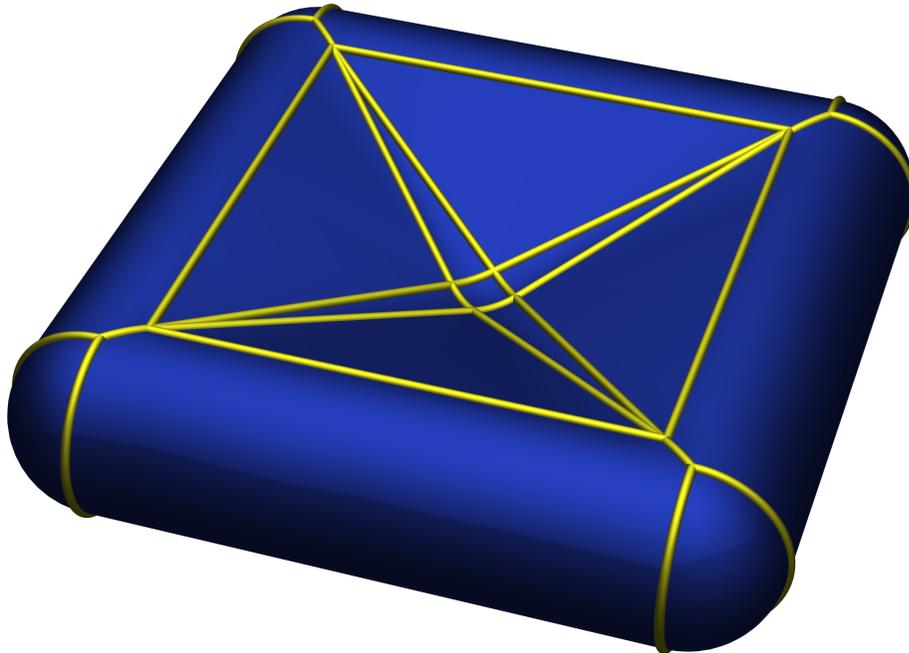


FIGURE 3.39. This hull consists of patches from round cylinders, cones, planes, and spheres. We outline the patches in yellow. In contrast to the tight hull boundary in Figure 3.37, the boundary from this hull does not contain patches from elliptical cylinders.

rounded square frame. We then add a square to the bottom of the frame and four triangles to its top. Each triangle is tangent at its base to one of the four circular cylinders in the frame and tangent at its tip to the green ball. Four conical patches complete the surface. The cones are tangent to the green ball and their apices are located where adjacent triangles touch.

This hull is not tight. We demonstrate this by appealing to a generalization of Gaussian curvature (Definition 10.17,) that is suitable for nonsmooth surfaces. According to the Gauss-Bonnet theorem [35, 59], the integral of Gaussian curvature over a watertight surface is determined solely by its topological type. The portions the hull's boundary that have positive Gaussian curvature are either locally convex or concave, while the negative Gaussian curvature is concentrated on the elliptic arc near each red ball.

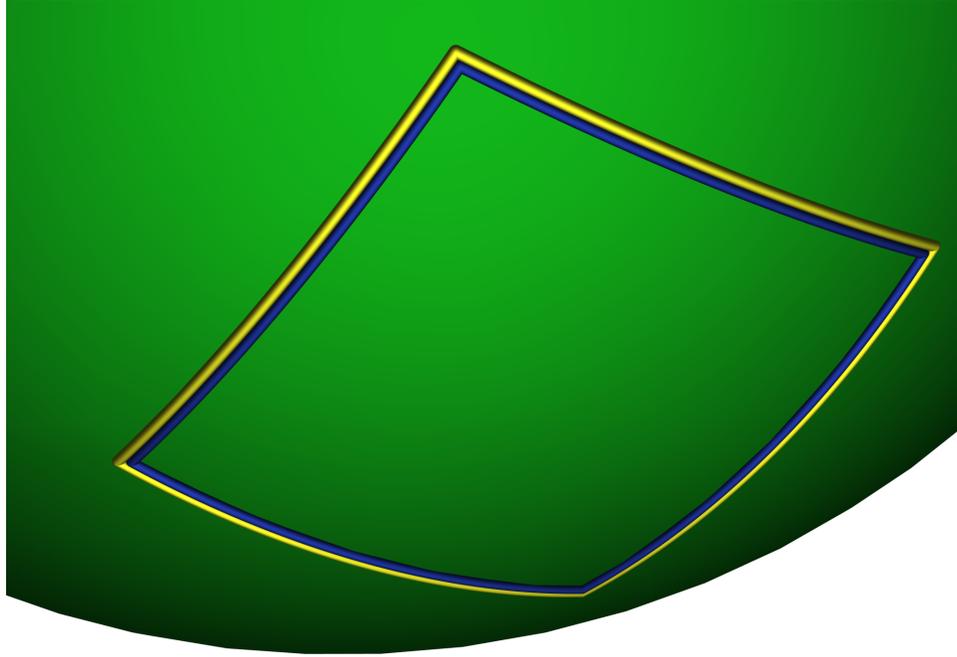


FIGURE 3.40. We show the green ball with the patch it contributes to the tight hull outlined in blue and the patch it contributes to the incorrect hull outlined in yellow. Although the patches are similar, the patch in the incorrect hull is slightly larger.

Let us denote the integral of the Gaussian curvature over the portion of the hull where the Gaussian curvature is positive by a , and the integral over the portion where the Gaussian curvature is negative by b . The total Gaussian curvature is $a + b$, and by the Gauss-Bonnet theorem $a + b = 4\pi$, which is the integral of Gaussian curvature over a sphere. The total absolute Gaussian curvature, which is equal to slack when we use the generalized form of Gaussian curvature, is $|a| + |b|$. Combining these results, we conclude that a hull is tight if and only if it minimizes a , which is necessarily nonnegative.

The subsets of the correct and incorrect tight hulls that have positive Gaussian curvature consist solely of spherical caps. The convex spherical caps in the two surfaces are identical, but each surface also contains a concave cap that lies on the boundary of the green ball, and the two concave caps are different. The cap in the incorrect hull is larger (Figure 3.40,) and because spheres have constant Gaussian

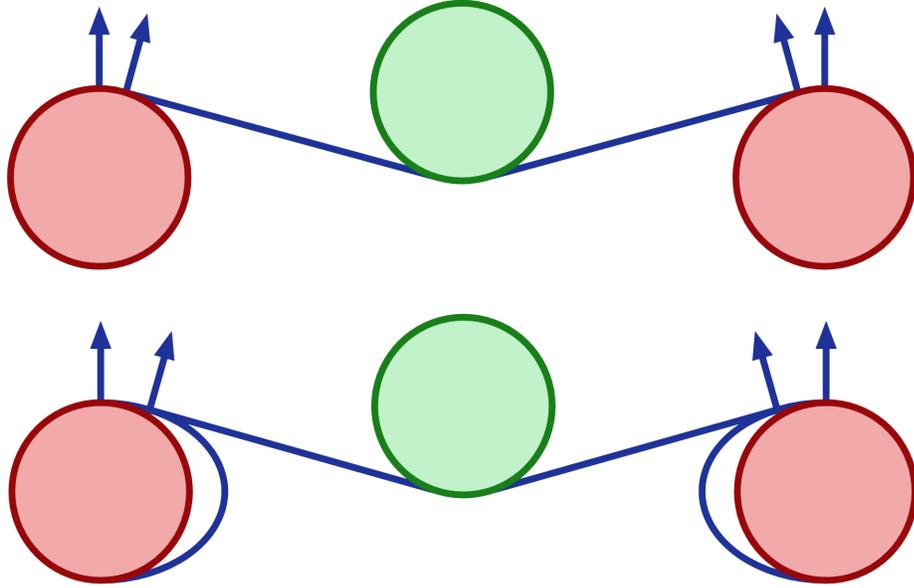


FIGURE 3.41. At the top, we show a section of the tight hull through the centers of the green sphere and two opposing red spheres. The two blue arrows indicate the change in normal direction over the circular arc lying on the edge of the elliptical torus patch in the tight hull boundary. We show the same section of the incorrect hull below. The arrows indicate the change in normal direction over the elliptic arc lying on the edge of the round torus patch in the incorrect hull boundary. The change in direction over the circular arc is less than the change over the elliptic arc. Accordingly, the change in angle over the arc on the boundary of the green disk is also smaller in the tight hull, corresponding to the smaller area of the patch in the tight hull from the green ball (Figure 3.40.)

curvature, we conclude that the value of a is larger in the incorrect hull, so it does not minimize slack.

The negative Gaussian curvature in these hulls is concentrated on their singular arcs, and we observe that the angle change over the circular arc, which lies on the elliptic cylinder in the correct tight hull, is smaller than the angle change over the elliptic arc, which lies on the circular cylinder in incorrect hull (Fig 3.41.)

If we rotate this section through the red and green balls about a vertical axis through the green ball, we obtain a solid torus R' and a green ball G' . The tight hull of R' relative to G' consists of a disk and patches from a torus, a cone, and a

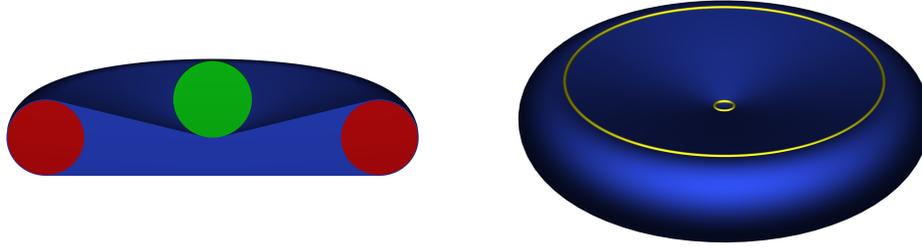


FIGURE 3.42. On the left, we show a cutaway view of the tight hull of the red torus R' relative to the green ball G' . In the image, R' is red, G' is green, and the hull is blue. On the right, we outline the toroidal, conical, and spherical patches in the hull boundary in yellow. There is an additional patch in the shape of a disk on the hull's bottom.

sphere (Figure 3.42.) This set has the same slack as the tight hull of R relative to G , but less of its slack is supported on R and G . We can envision a slack-preserving deformation that bends the toroidal sections in the tight hull of R' relative to G' connecting adjacent spheres in R into the circular and elliptic cylinders in the tight hull of R relative to G . The deformation transfers positive Gaussian curvature on the outside of the torus onto the spherical caps in the tight hull of R relative to G , and it concentrates negative Gaussian curvature from the inside of the torus on the correct hull's singular arcs.

TIGHTENING OVERVIEW

4.1. Tightening bounds curvature

The two-dimensional tight hull of R relative to G is guaranteed to be smooth when the boundaries of R and G are smooth. This is not true in higher dimensions, as the solution to the tablecloth problem illustrates. In two dimensions, all hull-like sets (as formalized by the hull property in Definition 12.2) have rubber band-like boundaries that press against the convex portions of the boundaries of R and G (Figure 4.1; Section 12.3.) If we scale each inward-pointing normal at each point on the boundary of the tight hull of R relative to G by the signed curvature at the point, the nonzero vectors point toward the interiors of R and G , provided R and G have nonempty interiors and smooth, manifold boundaries (Figure 4.2.) The boundary of

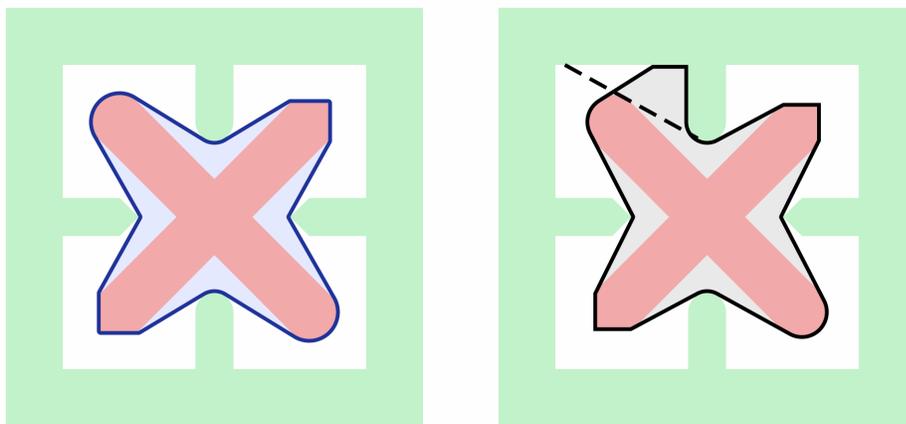


FIGURE 4.1. The dark blue hull boundary on the right intersects the boundaries of R and G , shaded red and green, only where they are convex. The boundary of the set shaded gray on the left, which intersects the boundary of G where it is concave, can be shortened by shaving the set at the dashed line. The boundary does not behave like a rubber band in tension, so the gray set is neither a tight hull nor a relative convex hull.

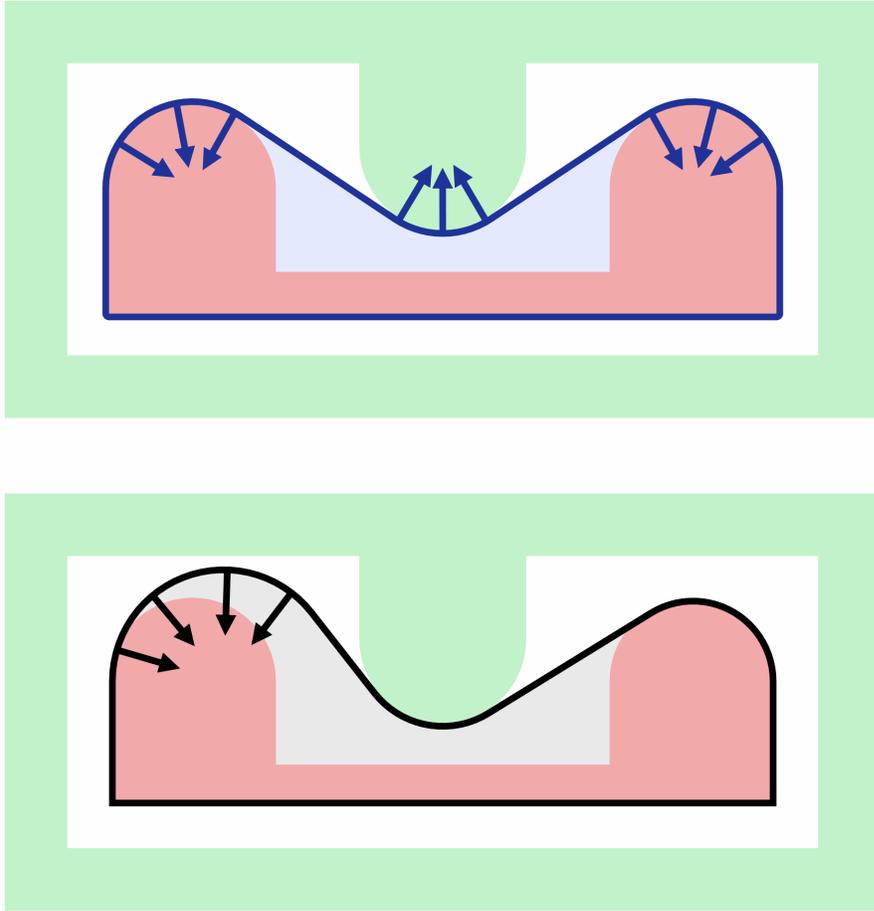


FIGURE 4.2. Above we show a hull boundary with the inward-pointing normals at several points scaled by their corresponding signed curvatures. Each scaled normal points toward the interior of either R or G , so the hull's boundary is invariant under curvature flow. At the bottom, we show a set in gray that is not invariant under curvature flow. A portion of its boundary is both curved and disjoint from the boundaries of R and G . That portion of the boundary would be displaced by curvature flow.

two-dimensional tight hull is then invariant under constrained curvature flow, so it locally minimizes length [26, 68].

Because it is invariant under curvature flow, the unsupported subset of the boundary of a two-dimensional hull has zero curvature. We show in Section 12.3 that it consists of line segments tangent to the boundaries of R and G . As a result, the normal variation over an open interval of the boundary of a two-dimensional hull is no greater than the normal variation of the interval's support (Figure 4.3.) This implies

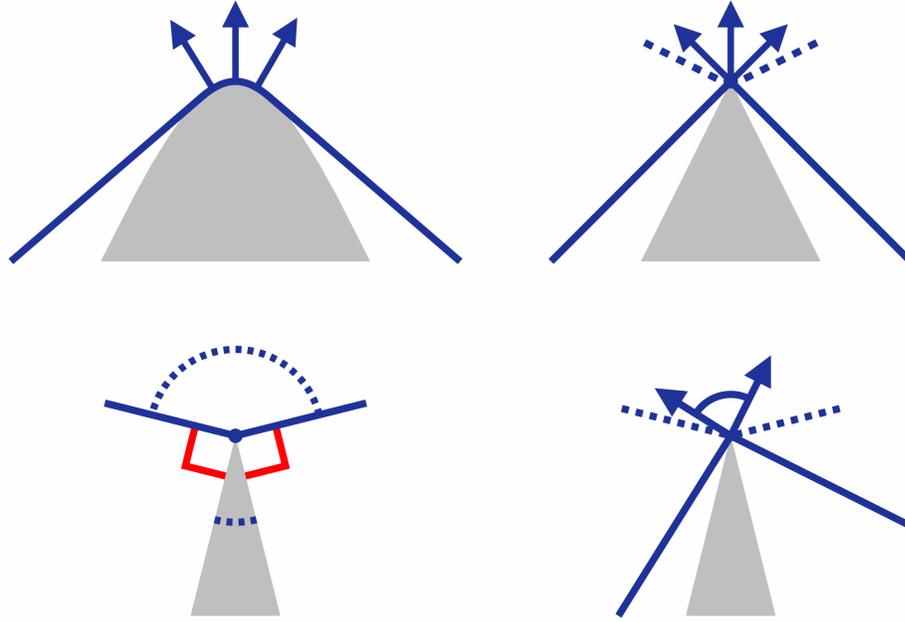


FIGURE 4.3. On the upper left, the hull boundary is smooth, and its normal variation precisely equals the normal variation of the supporting set in gray. On the upper right, we use blue arrows to indicate the fan of normals to the hull boundary at the marked vertex, and we use two dashed blue segments to indicate the fan of normals to the gray support. The normals to the hull are a subset of the normals to the support, and they are the only contributor to the hull's normal variation over any small interval containing the vertex. At the lower left, we illustrate how the cone of normals at a vertex on a polygonal curve subtends an arc equal to the difference between π and the internal angle between the edges incident on it. In the figure, the red brackets indicate two right angles subtending a total of π radians, so the two dashed arcs also subtend a total of π radians. The interval of angles spanned by the normal fan at the vertex is then the difference between π and the angle between the edges incident on the vertex. At a vertex of a polygonal set's hull, the interval of directions between the bounding edges of the constraint set incident on the vertex is always contained in the interval of directions between the hull edges incident on the vertex. As we illustrate on the lower right, this implies the fan of normals to the hull is always contained within the fan of normals to the constraints.

that the curvature of the boundary of a tight hull of R relative to G is no greater than the maximum curvature of the locally convex subsets of the boundaries of R and G , which motivates our definition of an r -tightening:

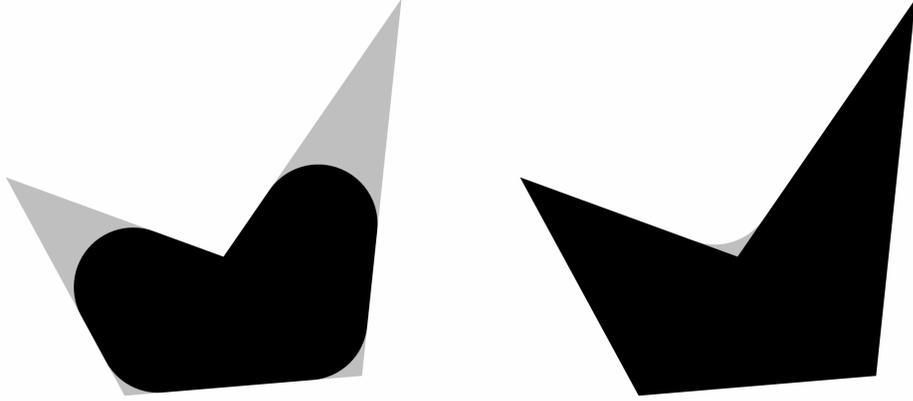


FIGURE 4.4. On the left, we show a polygon in gray with its morphological opening by a disk superimposed in black. On the right, we show the polygon’s morphological closing by the same disk in gray, with the polygon superimposed in black.

Definition 4.1. For $S \subseteq \mathbb{R}^d$, the *morphological opening* of S with radius r , denoted $S \circ_r$, is the union of all d -dimensional open balls of radius r contained in S . The *morphological closing* of S with radius r , denoted $S \bullet_r$, is defined as $S^c \circ_r^c$ (Figure 4.4.)

An *r-tightening* of S is a tight hull of $S \circ_r$ relative to $S^c \circ_r$ (Figure 4.5.)

The boundary of a tightening lies in the mortar, which is the set difference $S \bullet_r \setminus S \circ_r$ between the closing and opening of S . We can visualize the mortar in terms of rolling a ball of on both sides the boundary of a set (Figure 4.6.) The *r-mortar* is the set of points that a ball of radius r cannot reach.

Because the the morphological r -opening of a set is a union of balls of radius r , a ball of radius r inside the opening touches every point on the opening’s boundary. If a point on the opening’s boundary is locally convex, the curvature at the point must be less than or equal to $\frac{1}{r}$, or the disk and opening’s boundary would intersect (Figure 4.7.) Because a tightening is a tight hull, its normal variation is no greater than than the normal variation of the locally convex subset of the boundary of its constraints. The locally convex subsets of the boundaries of the r -opening of a set S and the r -opening of the complement of S both have curvature magnitudes less than or equal

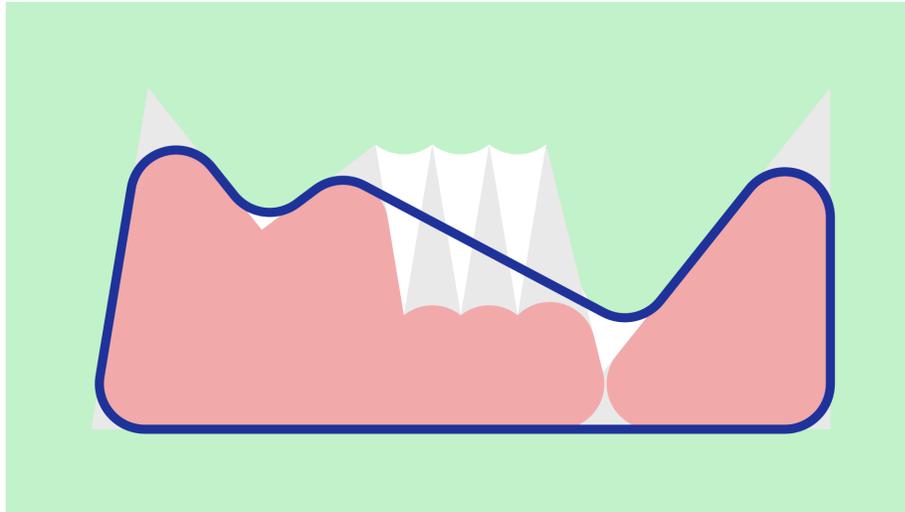


FIGURE 4.5. We show the opening of a polygon in red and the opening of its complement in green, with the difference between the polygon and its opening shaded gray. We outline the boundary of the polygon's tightening in blue. The blue curve bounds the tight hull of the red set relative to the green set.

to $\frac{1}{r}$, so the magnitude of the curvature along the boundary of an r -tightening is also less than or equal to $\frac{1}{r}$ (Chapter 14.)

Under the definition of curvature as the rate of change of the normal as a function of arclength, the curvature is only defined where the normal field is differentiable. A tightening may contain points where the normals are continuous but not differentiable (Figure 4.8,) so in Section 10.5 we provide an alternative definition of curvature that is applicable wherever the normal field is continuous. With this definition, we can assert that every point on the boundary of a tightening hull has a curvature in $[-\frac{1}{r}, \frac{1}{r}]$.

4.2. Continuously tightening polygons

We consider tightening polygonal data to be of special interest. In part, this because polygonal shape representations are ubiquitous, and in part it is because the tightening of a polygon has a simple structure. It consists of an alternating sequence of line segments and circular arcs, where each segment is tangent to two arcs. In singular situations, a segment may have zero length (Figure 4.9.)

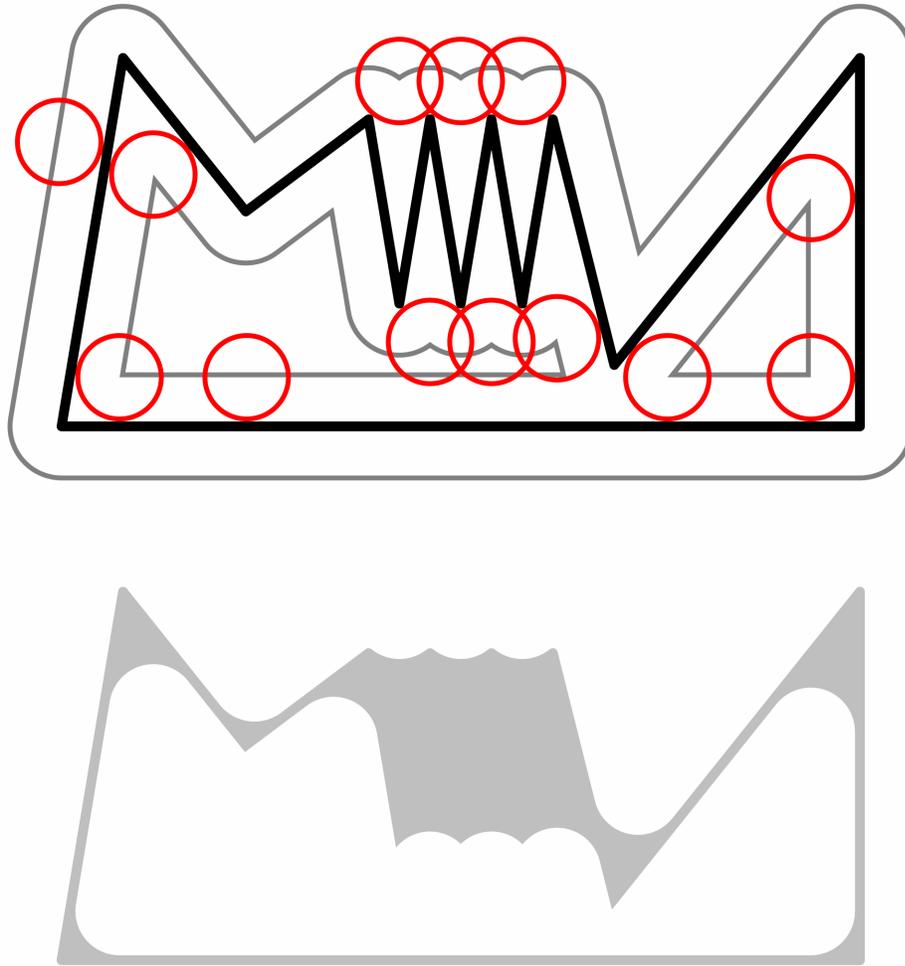


FIGURE 4.6. Above, we outline a polygon in black, while in gray we show the trajectories of two balls of radius r rolling along its boundary from both its inside and outside. We outline the rolling ball in red at several positions along the gray trajectories. Below, we show the polygon's r -mortar, which is the set of points that cannot be contained in a ball of radius r lying completely within it or its complement. The mortar collapses to a curve along some of the polygon's edges, but its interior is nonempty where the polygon has a sequence of spikes. We thicken the mortar slightly in this figure to make its lower-dimensional portions visible.

For every polygon with a manifold boundary, there is a value of r small enough so that the polygon's r -tightening replaces each of its vertices with a circular arc. We interpret the arc as performing a linear interpolation between the normals of its tangent edges. While the change in normal occurs as an impulse at a vertex,

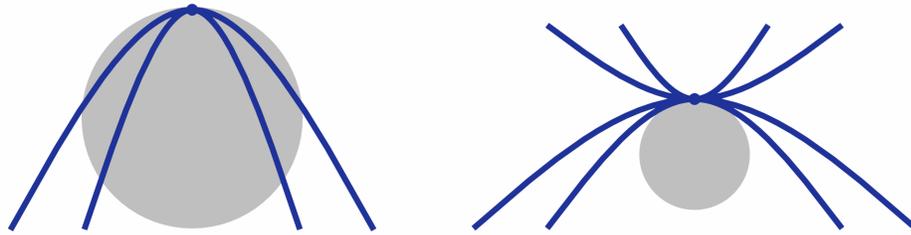


FIGURE 4.7. On the left, we show two curves in blue and a gray disk of radius r passing through the same point with the same normal. Both of the curves have a curvature greater than $\frac{1}{r}$ at the point indicated by the blue dot, and so they locally lie within the disk. On the right, we show a similar figure with four curves and a gray disk. All of the curves have a curvature less than $\frac{1}{r}$ at the marked point (the curvature of the two upper curves is negative,) so none of them intersect the interior of the gray disk in a small neighborhood of the point.

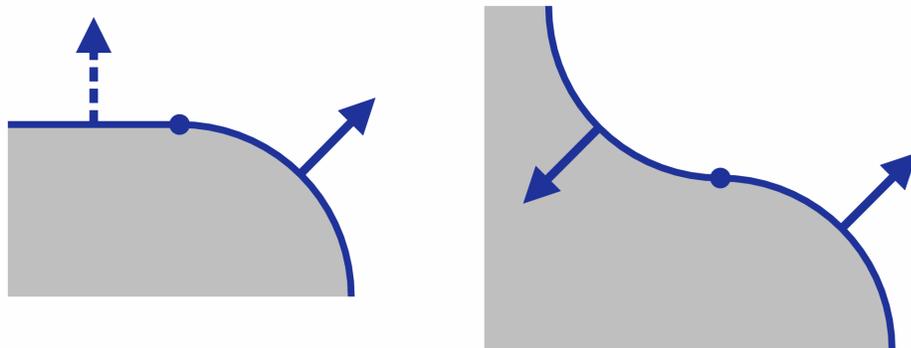


FIGURE 4.8. On the left, we show a line segment tangent to an arc from a circle of radius r at a point indicated by a blue dot on the boundary of a set shown in gray. Although the normals to the set's boundary are continuous, they are not differentiable. The normals' rate of change is zero over the line segment (which we indicate by drawing the normal with a dashed line) and equal to $\frac{1}{r}$ over the arc. On the right, we show a similar figure where two arcs from circles of radius r meet. Although the field of outward-pointing normals to the gray set is continuous at the marked point, the rate of change of the normals switches from $-\frac{1}{r}$ over the left arc to $\frac{1}{r}$ over the right arc.

the arc uniformly distributes the change in the normal over a curve whose length is proportional to r (Figure 4.10.)

As we continuously increase the value of r , the lengths of the arcs increase and the rate of change of the normals across them decreases. Meanwhile, the lengths of

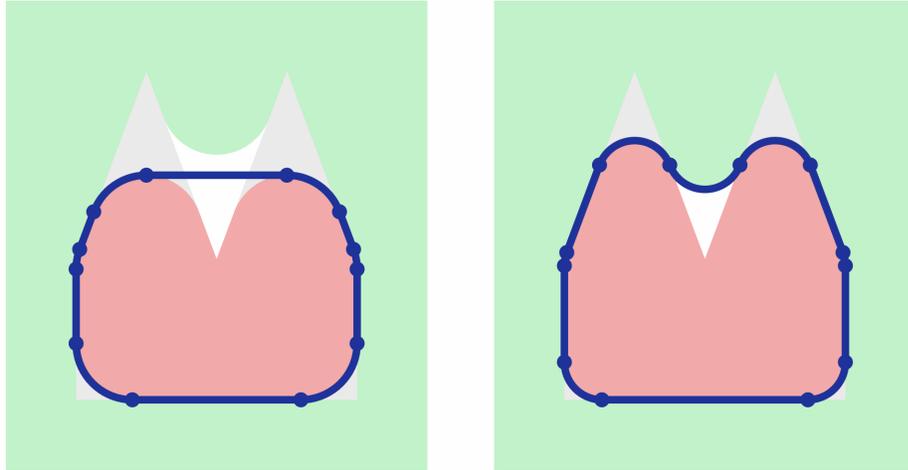


FIGURE 4.9. The r -tightening of a polygon typically consists of an alternating sequence of line segments and arcs from circles of radius r , as shown on the left. Each segment is tangent to two arcs, one at each endpoint. On the right we show a singular situation where two arcs meet with normal continuity. We can consider there to be an edge of length zero connecting the arcs.

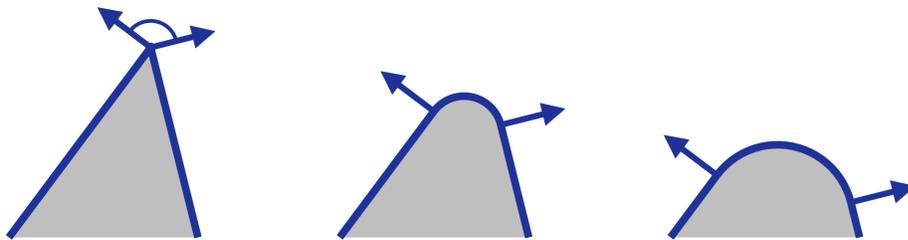


FIGURE 4.10. On the left, we show the change in normal direction that occurs at a vertex as we move from one edge incident on the vertex to the other. In the middle, we replace the vertex with a small arc, distributing the change in direction over the arc. On the right, we replace the vertex with a large arc. The change in normal direction occurs more slowly as we traverse this arc than as we traverse the small arc in the middle or pass over the vertex on the left.

the edges that connect the arcs decrease (Figure 4.11, left.) At some value of r , two arcs meet and an edge disappears (Figure 4.11, center.) As r continues to increase, a new edge appears (Figure 4.11, right.) Unlike the other edges, which lie on the polygon's boundary, the new edge is disjoint from the opening of the polygon and its complement, making it unsupported.

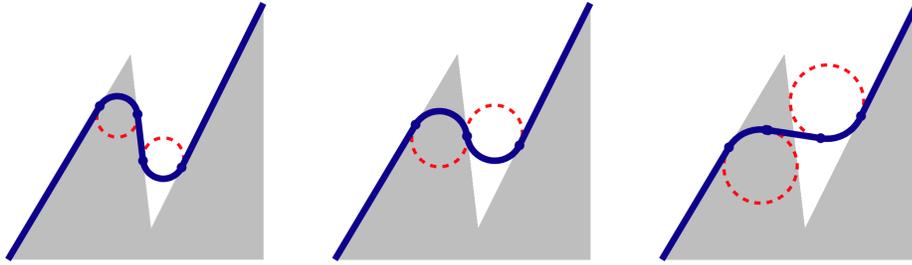


FIGURE 4.11. On the right, we show a portion of a gray polygon tightened with a small radius. We draw the boundary of the tightening in blue, with dots marking the endpoints of two arcs and the edge connecting them. In the middle, we increase the radius used for tightening. This makes the arcs longer, and their endpoints slide toward each other along the polygon's edge, eliminating the edge that connected them on the right. On the left, we show how further increasing the tightening radius introduces an edge into the tightening's boundary that does not lie on an edge from the input.

Morphologically opening a set may change its topology. Accordingly, there may be a value of r where the topology of a set's r -tightening changes, producing a discontinuous change in its boundary (Figure 4.12) Over intervals of r free of topological changes, tightening defines an isotopy. We show in Chapter 14 that for a large class of sets, the number of topological changes is finite, so that we can characterize tightening as piecewise continuous.

4.3. Tightening simplifies normal fields

Over an interval of values of r where the boundary of an r -tightening is continuous, its slack monotonically decreases. The set X of all sets that contain S_{\circ_r} and exclude $S^c_{\circ_r}$ grows larger with r because the morphological opening of a set with radius r shrinks as r grows; the opening of a set with radius $a \in \mathbb{R}^>$ is a subset of its opening with radius $b \in \mathbb{R}^>$ if $a > b$. As the value of r increases and X grows larger, the minimum slack over the sets in X can only decrease.

Over an interval where the tightening is continuous, we link decreases in slack to the behavior of unsupported edges. These edges act like shortcuts that trim slack and unfold a curve we call the Gauss offset.

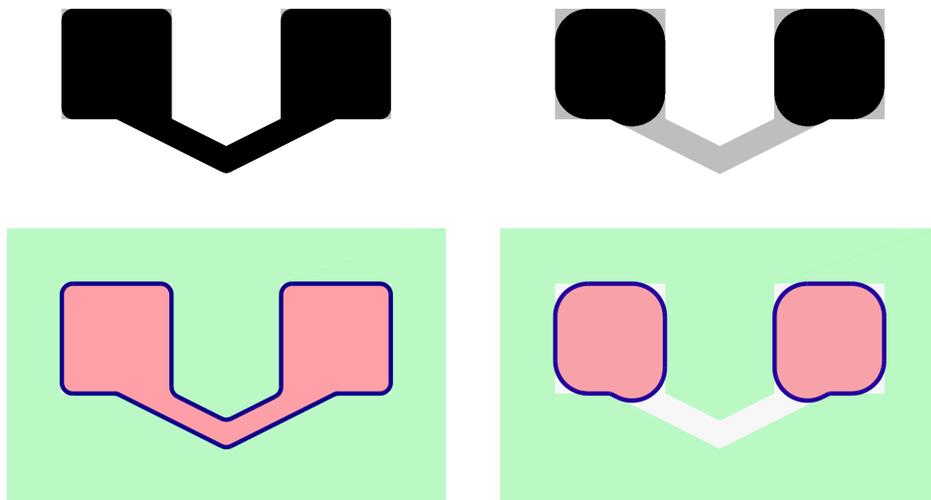


FIGURE 4.12. On the top left, we show a polygon in gray with its morphological opening by a small disk superimposed in black. On the top right, we show the same polygon with a larger-radius opening superimposed. On the bottom left, we show the boundary of the tightening of the polygon with a small radius in dark blue. We shade the polygon's opening red and the opening of its complement green. On the bottom right, we show a similar figure illustrating tightening with a larger radius. Although the tightening on the bottom left has one connected component, the tightening on the bottom right has two.

To define the Gauss offset of a set $S \subseteq \mathbb{R}^d$, we need S to have a well-defined normal field. A tightening's normal field is given by the outward-pointing normals of its bounding arcs and edges. Because a tightening contains no cusps or corners, all of its points have unique normals, including those points where an arc and an edge or two arcs meet. A polygon has a fan of normals at each vertex swept by the normal to one of its incident edges as we rotate it to align with the normal to the other edge (Figure 4.13). In Section 10.3, we provide a more general definition of a normal field that is applicable to arbitrary subsets of Euclidean space.

The ϵ -Gauss offset of $S \subseteq \mathbb{R}^d$ is the result of scaling S by a small factor $\epsilon \in \mathbb{R}^>$ and displacing each point on the boundary of S by its unit normals. If S is planar, the ϵ -Gauss offset of S converges to a circle as ϵ approaches zero, but it has folds if S is nonconvex (Figure 4.14.)

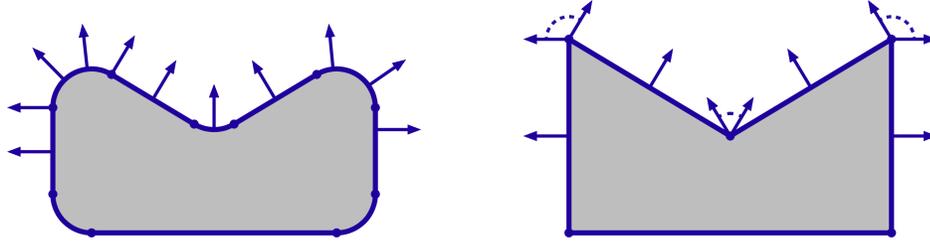


FIGURE 4.13. The normal field for the set produced by tightening on the left is defined by the normals to its bounding arcs and edges. Because the edges of the polygon on the right meet at singular points, each polygon vertex possesses a set of normals covering the arc swept when we rotate the normal to one of its incident edges until it aligns with the normal to the other edge.

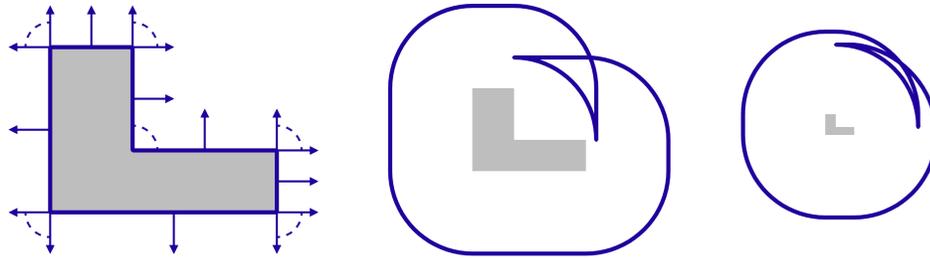


FIGURE 4.14. On the left, we show a nonconvex polygon in gray with its normal field indicated in blue. In the middle, we show the polygon's $\frac{1}{2}$ -Gauss offset, which self-intersects due to the polygon's concave vertex. On the right, we show the polygon's ϵ -Gauss offset as ϵ goes to zero. The offset approximates a circle, but it has folds with patches where geodesically distant points with similar normals map to nearby points under the Gauss offset.

Suppose we define a function of direction that is equal to the number of times a ray from the circle's center with a given direction intersects the Gauss offset. As ϵ goes to zero, the value of the function for a direction converges to the number of connected components from the set's boundary whose normals have that direction (Figure 4.15) We refer to this function as the polygon's normal count (Definition 10.12.)

If S is a polygon, an edge of length l in S maps to an edge of length $l\epsilon$ under the ϵ -Gauss offset of S . As ϵ goes to zero, the contribution from the edges of S to the length of the ϵ -Gauss offset goes to zero. The length of the Gauss offset converges to the integral of the normal count over the circle of directions, so the edges of S do not contribute to this integral. A vertex, by contrast, maps to a circular arc in the

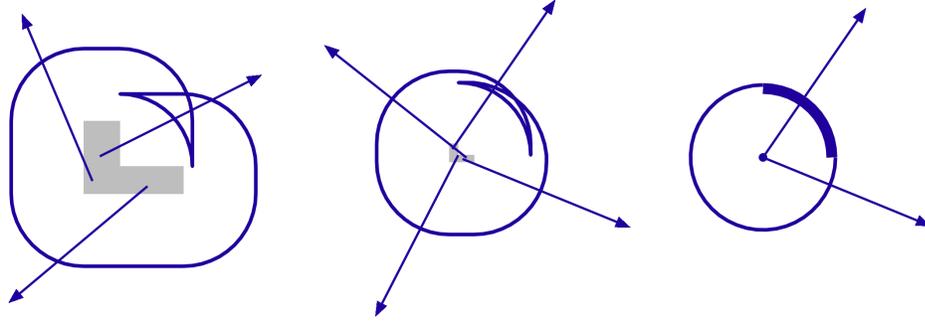


FIGURE 4.15. If we pick a ray with its tail at a point in a set and count the number of times the ray intersects the set's ϵ -Gauss offset, the result depends on the tail position and on the precise value of ϵ when ϵ is large. As ϵ decreases, the set of tail positions converge to a point and the geometry of the Gauss offset stabilizes. As ϵ approaches zero, the mapping from ray directions to the number of ray intersections with the ϵ -Gauss offset converges, defining a function we call the set's normal count.

Gauss offset, and it adds one to the normal count over an interval of directions. The integral of the normal count and the length of the Gauss offset in the limit are both equal to the sum of the absolute value of the angle change at each vertex, which is the polygon's slack.

The normal count of the polygon is invariant during tightening as long as tightening only replaces each polygon vertex with an arc. The vertex and the arc replacing it contribute identical amounts of slack, as suggested previously in Figure 4.10. Every boundary segment connecting two points has the same slack, provided the change in normal direction over the segment is monotonic (Figure 4.16.) When an edge from the polygon disappears and an unsupported edge replaces it, the unsupported edge trims its incident arcs and reduces the total change in normal that occurs over them (Figure 4.17.) This corresponds to an unfolding of the Gauss offset and a reduction in normal count in the direction of the normal to the unsupported edge.

We formalize this result in Chapter 14. We combine tightening's slack-reducing and curvature-limiting properties to conclude that tightening a set simplifies its normal field. Reducing slack reduces the variation of the normals to a set, while limiting curvature disperses the set's normal variation over its boundary.

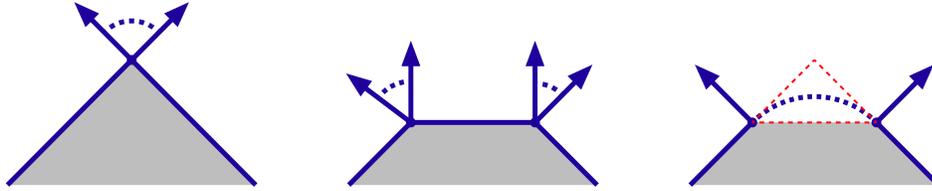


FIGURE 4.16. On the left, we show a corner of a polygon. Slack is concentrated at the corner and is equal to the angle between the two edges incident on the vertex. In the middle, we cut the corner, which divides the slack into two equal parts located at the shown vertices. We are given two edges with a gap between them on the right. Any convex function defined over the dashed red line that remains within the gray triangle has slack equal to the magnitude of the difference in direction between the edge normals. We indicate one such function with a dashed blue curve.

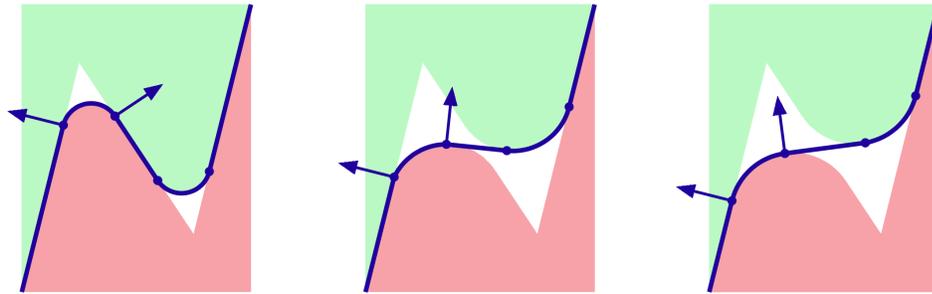


FIGURE 4.17. On the left, we show a portion of the r -tightening of a polygon. We mark the endpoints of two arcs from the tightening and the tangent edge connecting them with blue dots. Two arrows indicate the change in normal direction over one of the arcs. In the middle, we show the tightening of the same polygon with a larger value of r . The edge that coincided with the boundary of the opening of the polygon and its complement is now unsupported, and only portions of the two arcs remain on the boundary of the tightening. On the right, we further increase the value of r . The length of the unsupported edge increases while the lengths of the arcs incident on it decrease. The change in normal direction indicated by the angle between the two arrows is significantly smaller than on the left, corresponding to a reduction in slack.

CHAPTER 5

MEDIAL COVER OVERVIEW

Computing the tight hull of R relative to G is straightforward when R and G^c are simple polygons. The space between R and G is then a topological annulus. We can cut this annulus by subtracting a line segment from it. We select a vertex of R that necessarily lies on the boundary of its tight hull, such as its leftmost vertex v . We replicate v , yielding a vertex w at the same position as v . We then cut the annulus using a ray extending leftward from v , introducing two edges and two new vertices, which we link into a simple polygon. (More precisely, the polygon is weakly simple [62] because the two edges we introduce spatially coincide.) We then find the shortest path between vertex v and w in linear time using the funnel algorithm ([11, 39, 62], Figure 5.1.)

We are unaware of any previously-known algorithm for computing the tight hull or relative convex hull when R consists of multiple connected components. (Toussaint presents a worst-case optimal $O(n \log n)$ algorithm for the geodesic hull of a set of points within a polygon [62], but the fact that the geodesic hull is connected makes its construction substantially different from that of a tight or relative convex hull.) When R consists of multiple components, there are multiple sets that contain R , exclude G , and possess manifold, rubber band-like boundaries. We refer to these sets as tight covers (Section 12.4; Figure 5.2.) The number of tight covers may be exponential in the number of components in R (Figure 5.3, Section 12.3) and a tight hull algorithm must determine which components of R lie in each component of the tight hull. If G^c contains holes, the algorithm must also determine how the hull boundary winds around those holes (Figure 5.4.)

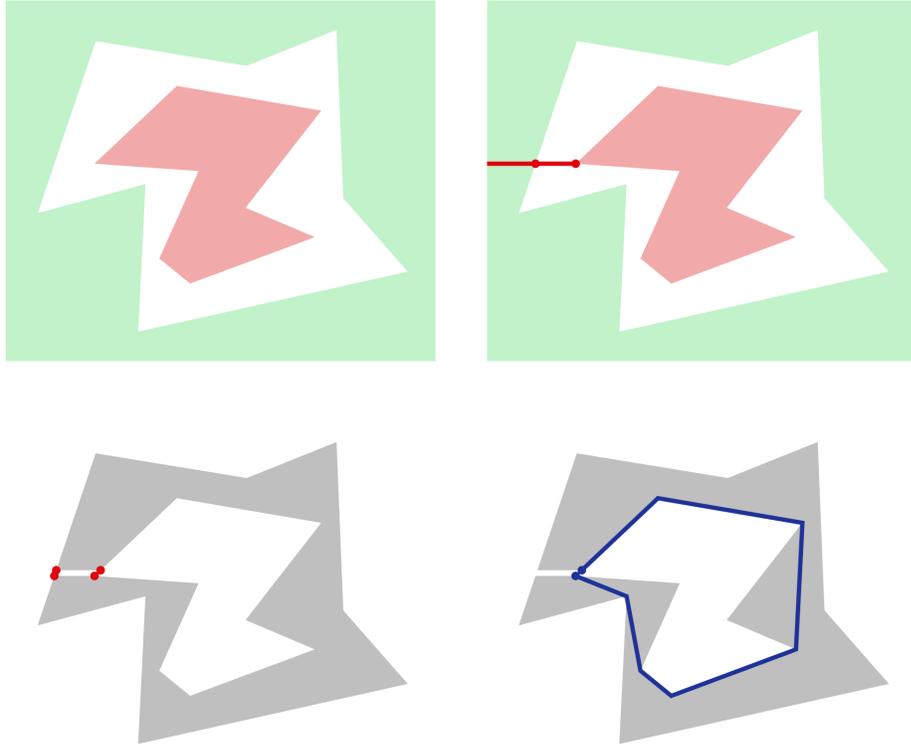


FIGURE 5.1. On the upper left, we show a set R in red and G in green, where R and G^c are simple polygons. The set $(R \cup G)^c$, which is the gap between R and G , is isotopic to an annulus. On the upper right, we cut the annulus horizontally with a ray extending leftward from the leftmost vertex v on the boundary of R , introducing a vertex p on the boundary of G . We show vertices v and p as well as the ray \vec{vp} in red. On the lower left, we symbolically duplicate the cut and both of its endpoints to obtain a simple gray polygon P with new vertices w and q . We show w and q in red above v and p , respectively. On the lower right, we construct the shortest path from v to w in linear time. Merging the path endpoints, we obtain the shortest loop in the annulus, which is the boundary of both the tight hull and convex hull of R relative to G .

We define the medial cover as an example of a specific tight cover that we can efficiently compute. In our definition (Chapter 15,) we use the medial axis to specify a set of points that we add to R and G (Figure 5.5.) These points function as a collection of barriers, rendering impossible any hull whose boundary crosses a barrier. Assuming that no four vertices from the the input are cocircular, the medial cover is the unique tight cover whose bounding loops are disjoint from the set of added

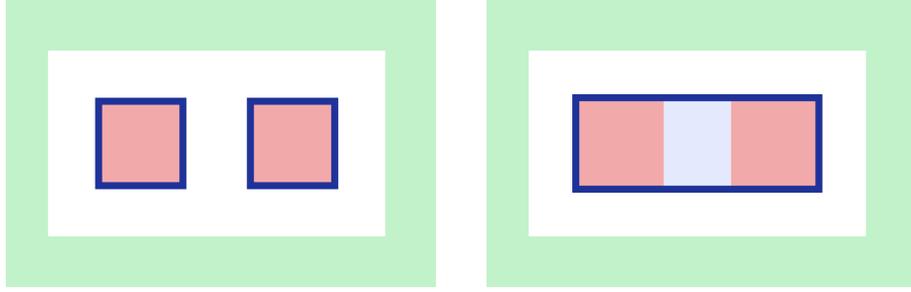


FIGURE 5.2. On the left, we outline in blue the boundary of a tight cover of a set R consisting of two squares relative to a set G equal to the complement of a rectangle. The cover is equal to R and has two connected components. On the right, we outline the boundary of a tight cover of R relative to G consisting of a single component. This tight cover is equal to the tight hull of R relative to G . We shade the points the cover adds to R light blue.

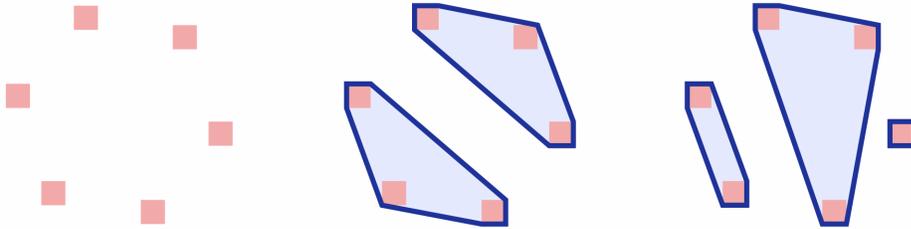


FIGURE 5.3. On the left, we show several small connected components of R arranged at points on a circle. Taking G to be empty, we show two tight covers of R relative to G at the middle and right. As shown in Section 12.3, the number of covers is not bounded by a polynomial function of components of R .

points. It can also be characterized as the tight cover whose boundary is isotopic to the set of points equidistant from R and G (Figure 5.6,) given that the deforming set contains R and excludes G throughout the isotopy.

Like the tight hull, the medial cover is symmetric with respect to set complement (Definition 8.28;) the set of points equidistant from R and G is the same as the set equidistant between G and R . The convex hull of R relative to G , by contrast, is asymmetric. It is the largest tight cover of R relative to G when the tight covers are ordered by set inclusion (Chapter 15,) implying that the complement of the convex hull of G relative to R is the smallest (Figure 5.7.)

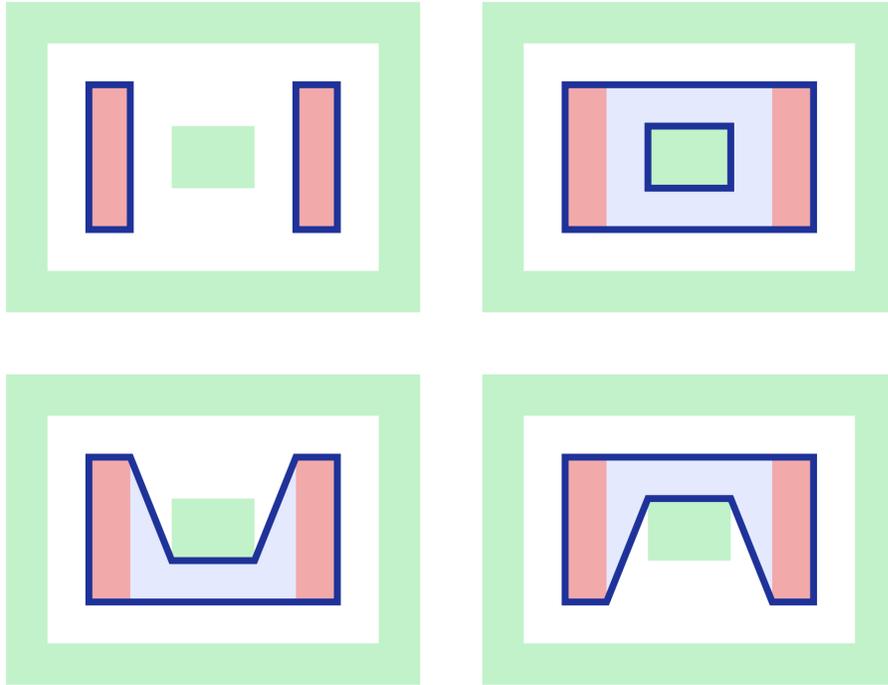


FIGURE 5.4. We show four tight covers of R relative to G . The cover in the upper left consists of two connected components. The other three covers each consist of a single component, although the component from the cover at the upper right has a hole. The two covers at the bottom are isotopic to each other in the plane, but they cannot be deformed into one another if we constrain the deforming set to contain R and exclude G .

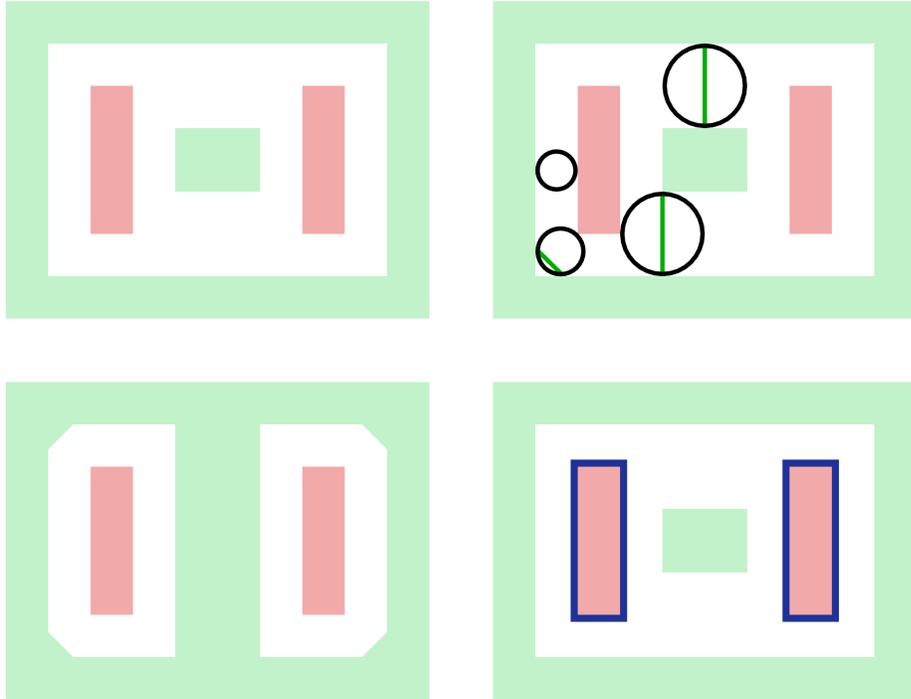


FIGURE 5.5. In the upper left, we show the sets R and G from Figure 5.4 in red and green, respectively. In the upper right, we show four of the maximal disks centered on the medial axis of $(R \cup G)^c$. For each disk, we take the convex hull of the points where it touches R and add the result to R . Similarly, we take the convex hull of the points where the disk touches G and add the result to G . In this example, we add no points to R , but we do add points to G . We show three of the line segments we add to G in green. In the lower left, we show R with the augmented green set G' . The set $(R \cup G')^c$ consists of disjoint annuli. In the lower right we show the medial tightening of R relative to G , which we obtain by constructing the shortest loop in each annulus. This medial tightening consists of two components. It is equal to the tight cover shown in the upper right of Figure 5.4, and it is different from the tight hull of R relative to G .

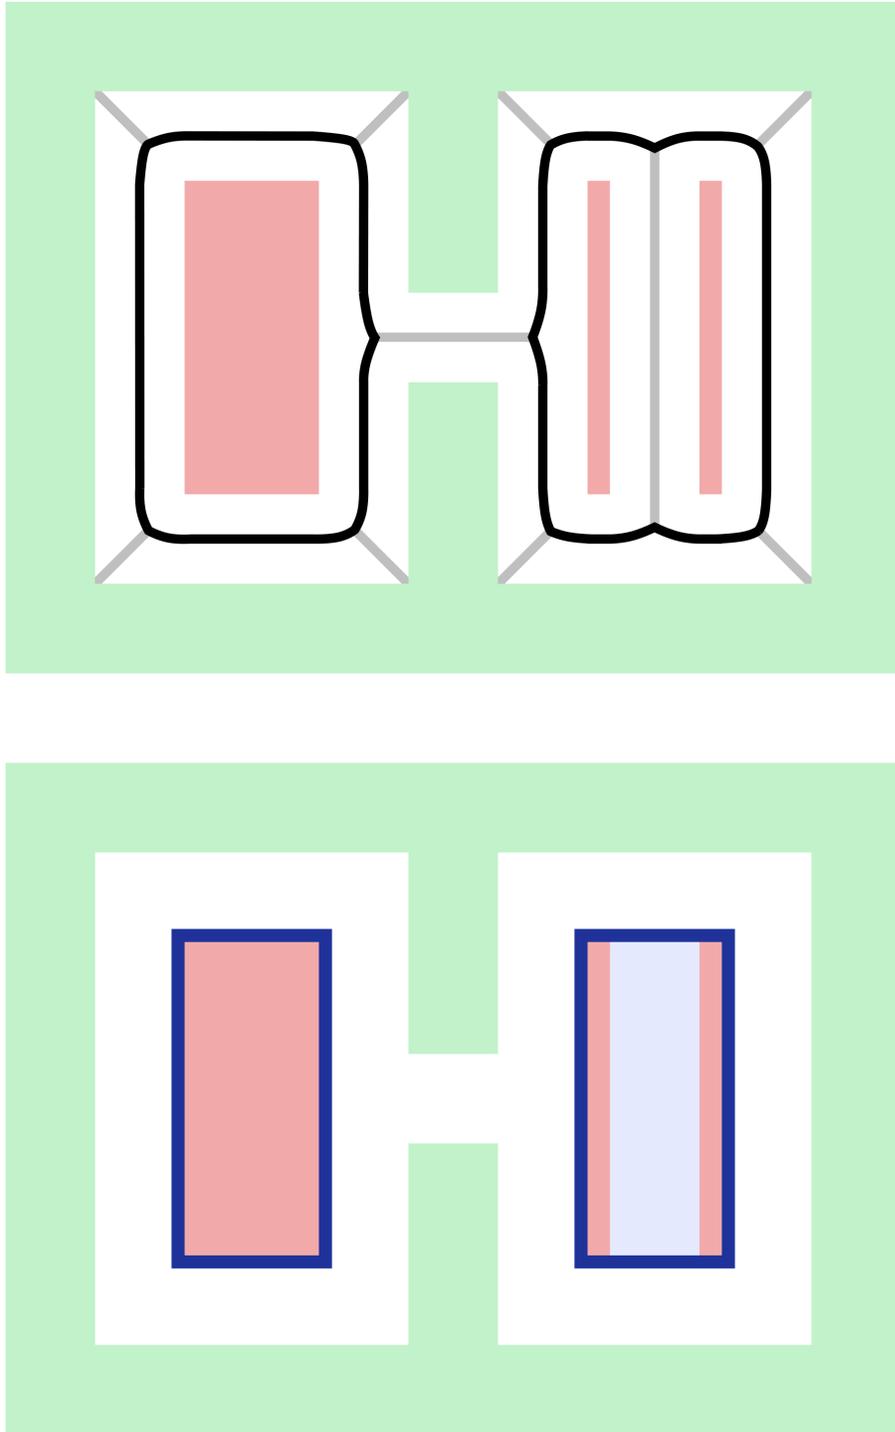


FIGURE 5.6. Above, we show R and G with the medial axis of $(R \cup G)^c$ in gray and the points equidistant from R and G (which form a subset of the the medial axis) in black. Below, we show the medial cover of R relative to G in blue. The black loops above may be deformed to the blue loops on the right without intersecting the interiors of R or G .

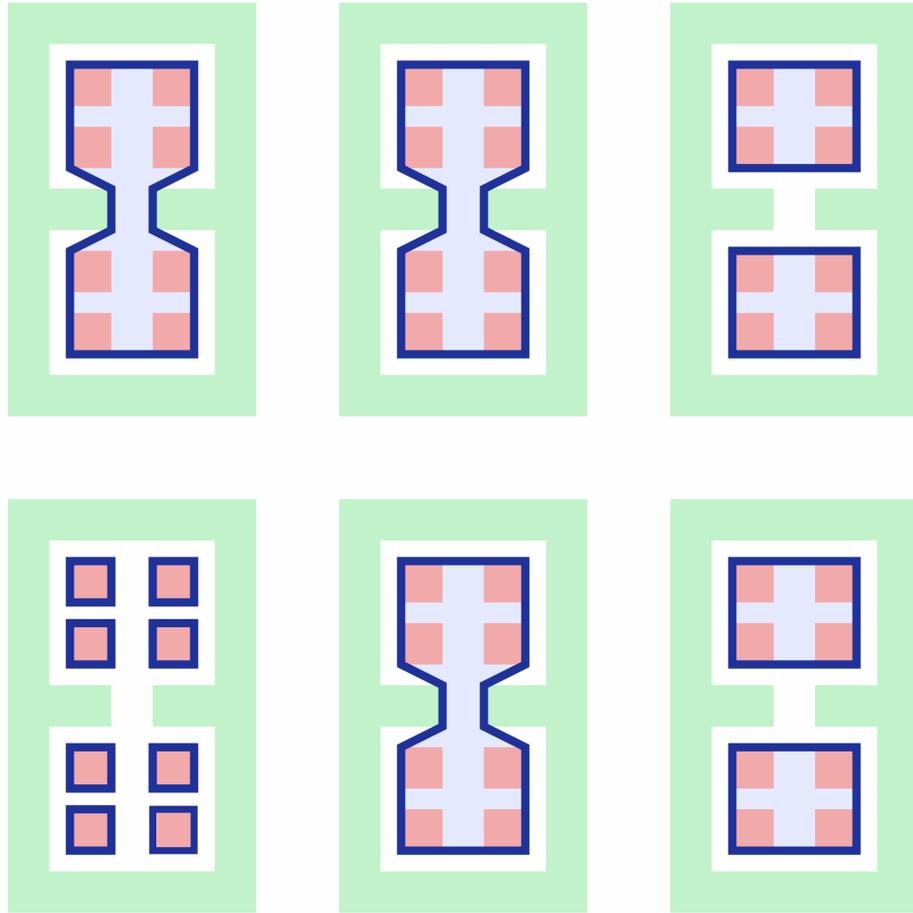


FIGURE 5.7. At the top left, we show the convex hull of R relative to G , which consists of a single large component. At the top middle, we show the tight hull of R relative to G , which in this case is identical to the convex hull of R relative to G . At the top right, we show the medial cover of R relative to G , which groups nearby components of R into two separate components. At the bottom left, we show the complement of the convex hull of G relative to R . While the convex hull of R relative to G is the largest tight cover, this cover is the smallest. At the bottom middle, we show the complement of the tight hull of G relative to R , and at the bottom right we show the complement of the medial cover of G relative to R . Because the tight hull and medial cover are symmetric with respect to set complement, the top middle and bottom middle images are identical, as are the images at top and bottom right.

CHAPTER 6

IMPLEMENTATION OVERVIEW

Assume that the complement of G is a simple polygon, and assume R consists of one or more disjoint simple polygons lying in G^c . We can compute a tight cover given any triangulation of the gap between R and G , and the triangulation requires time in $O(n \log n)$ if R and G have a total of n vertices. We define a triangle from the triangulation to be mixed if its vertices lie on the boundaries of both R and G . If its vertices lie only on R , it is red, and if they lie only on G , it is green (Figure 6.1.) Because each mixed triangle has exactly two mixed edges, each with one vertex on the boundary of R and the other on the boundary of G , a mixed triangle connects to exactly two other mixed triangles. As a result, the mixed triangles form loops or cycles that define annuli. The tight hull of the bounded component of the complement of each annulus relative to the unbounded component can be computed in linear time, as

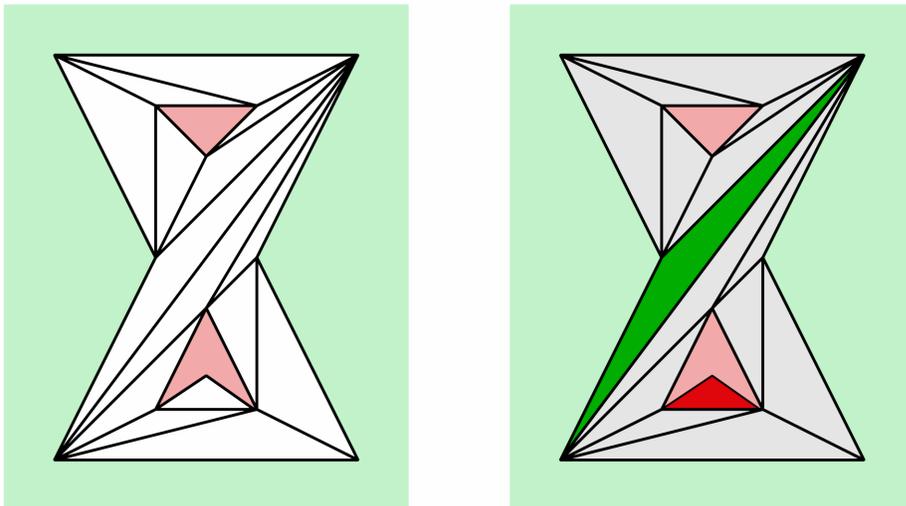


FIGURE 6.1. On the left, we show a triangulation of the space between R and G . On the right, we shade the triangles whose vertices lie on the boundaries of both R and G gray. We shade the triangle whose vertices lie on R red, and we shade the triangle whose vertices lie on G green.

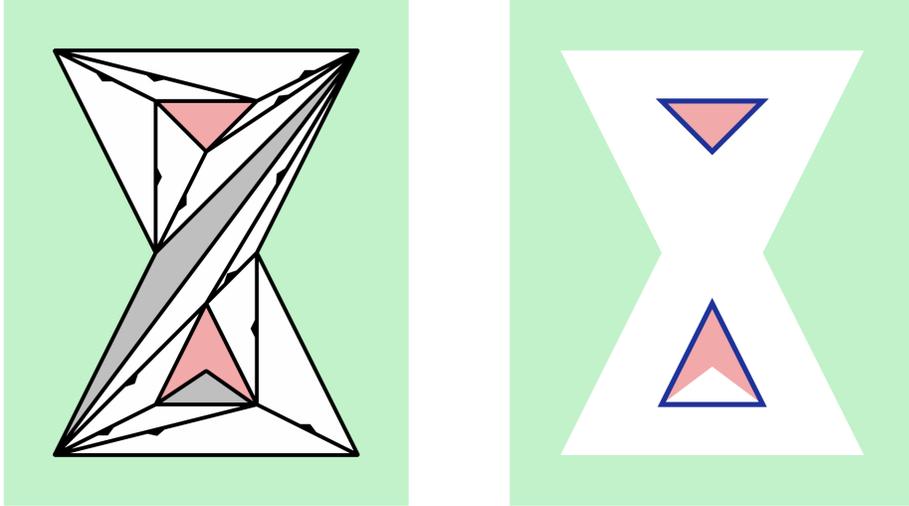


FIGURE 6.2. On the left, we illustrate how each mixed triangle connects to two mixed triangles, one across each of its two mixed edges. We shade the solid triangles gray. Using a consistent triangle orientation, we can assign each mixed triangle a predecessor and successor. We indicate this with a black triangle at the midpoint of each mixed edge that points from the mixed triangle preceding the edge to its successor. Because the number of mixed triangles is finite, they must form cycles, and the union of the triangles in a cycle forms an annulus. The unshaded triangles on the left form two annuli with disjoint interiors. On the right, we show the shortest loop in each annulus in blue. The cost to compute a loop is linear in the number of vertices in its annulus. The sum of the vertices in all the annuli is linear in the number of vertices in R and G , so the cost to compute all the loops is linear in the size of the input to the medial cover problem.

we explained in Chapter 5. The annuli decompose the medial cover problem into the problem of computing a linear number hulls of polygonal inputs with trivial topology. The decomposition keeps the total input size over all of the trivial hull subproblems linear in the medial cover input size. The total cost of path planning for all the annuli is then linear (Figure 6.2.) Consequently, we can compute a tight cover in $O(n \log n)$ time. The initial step of triangulating a polygon with holes dominates the complexity.

If we compute a tight cover from an arbitrary triangulation, the cover is also arbitrary. Because which cover an arbitrary triangulation produces is not tied to geometric properties of R and G , perturbing the geometry of R and G may unpredictably

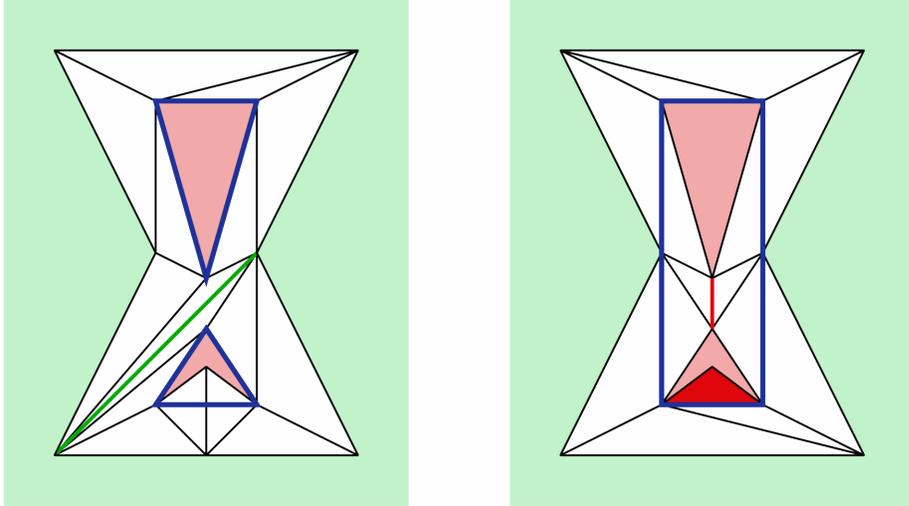


FIGURE 6.3. On the left, we show a tight cover of R relative to G constructed from an arbitrary triangulation. We shade the solid edge connecting two vertices of G green and outline the hull boundary in blue. On the right, we show the constrained Delaunay cover of R relative to G , which we construct from a constrained Delaunay triangulation of the space between R and G . We shade the solid edge and triangle red, with the boundary of the constrained Delaunay cover in blue.

change the cover, even if the triangulation algorithm is deterministic. This renders it impossible to compute an approximation of the piecewise continuous deformation caused by tightening with a progressively increasing radius.

We can address this limitation by using a geometrically-based triangulation, such as a constrained Delaunay triangulation. Given a representation of polygonal sets R and G , we define the corresponding constrained Delaunay cover as the cover determined by the annuli extracted from a constrained Delaunay triangulation of the input (Figure 6.3.) The constrained Delaunay triangulation can be computed in $O(n \log n)$ time with standard algorithms, but it is sensitive to the way the input representation samples the boundaries of R and G . If we insert vertices into a representation of the edges bounding R and G , the corresponding constrained Delaunay cover of R relative to G may change although R and G do not (Figure 6.4.) As a result, the constrained Delaunay cover is not exclusively a function of R and G as subsets of the plane. It depends on how we represent R and G as well.

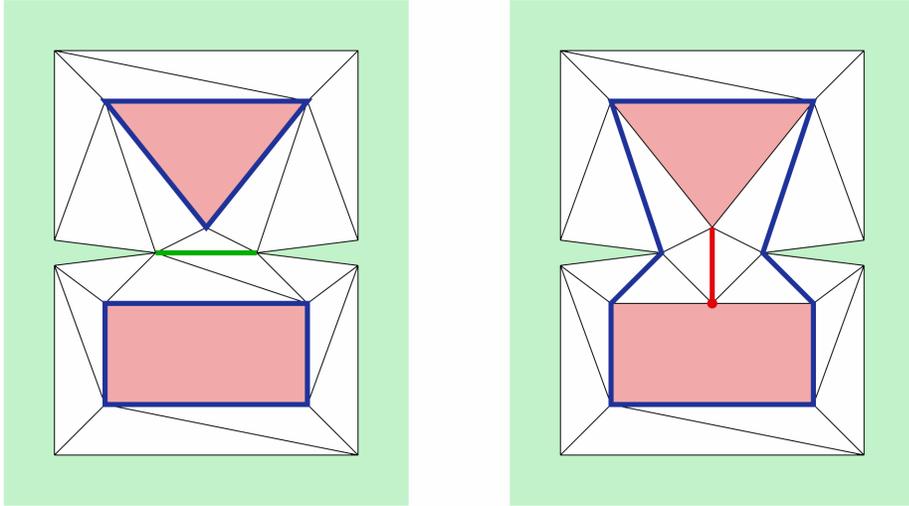


FIGURE 6.4. On the left, we show a constrained Delaunay cover constructed from vertices located at the corners of R and G . On the right, we introduce a vertex into one of the edges bounding R . The constrained Delaunay triangulation using that vertex is different from the one that does not, yielding a different constrained Delaunay cover.

If we densely sample the boundaries of R and G , the constrained Delaunay cover of R relative to G stabilizes as the spacing between vertices goes to zero. In a constrained Delaunay triangulation constructed from a dense sampling of the boundaries of R and G , there is a large number of thin triangles and a small number of large triangles. As we increase the sampling density, the number of thin triangles increases while the number of large triangles remains fixed (Figure 6.5.)

The circumcenters of the large triangles approximate bifurcation points of the medial axis (Figure 6.6.) We can obtain the bifurcation points by computing the entire medial axis of $(R \cup G)^c$, which requires $O(n \log n)$ time. We then compute the points where the maximal disk centered on each bifurcation point touches R and G , and we add those points to the input as vertices (Figure 6.7.) The constrained Delaunay cover of this modified input is the same as its medial cover; inserting additional vertices has no effect (Figure 6.8.) The total computation requires $O(n \log n)$ time. This is optimal in the worst case, because by making G^c large we can use an algorithm for the medial cover of R relative to G to compute the convex hull of R (Figure 6.9.)

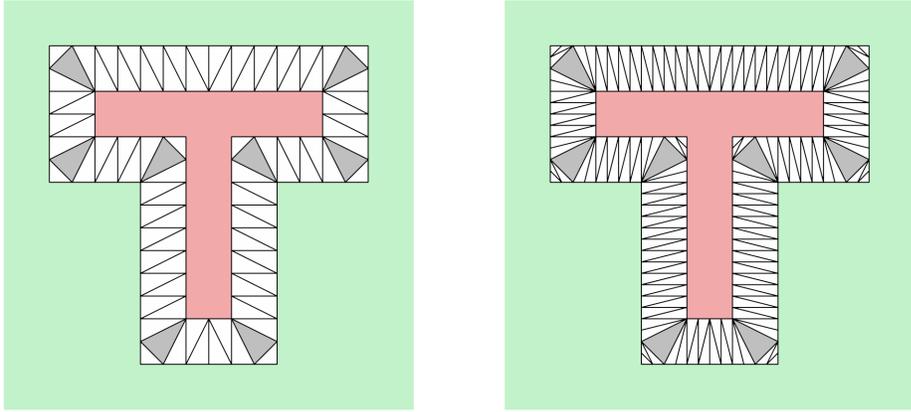


FIGURE 6.5. On the left, we show a constrained Delaunay triangulation of the space between uniformly sampled polygonal sets R and G . We shade gray the triangles that only touch the boundaries of R and G at their vertices. The gray triangles are located at corners of R and G , and they are slightly larger than the other triangles. On the right, we double the number of samples along the boundaries of R and G . This doubles the number of unshaded triangles and halves their widths. It does not affect the number or size of the gray triangles.

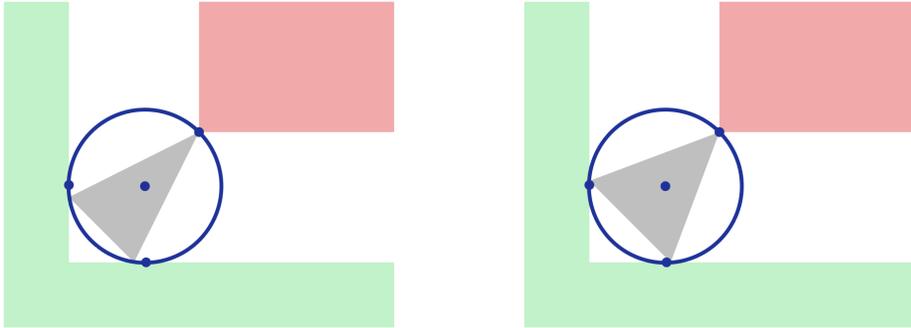


FIGURE 6.6. On the left, we zoom in on a portion of the sets R and G shown in Figure 6.5. The vertices of the large gray triangle are close to the points where a maximal disk centered on a bifurcation point of the medial axis touches the polygon. We outline the maximal disk in blue and mark its center and the points where it touches the boundaries of R and G with blue dots. At this scale, the gray triangle's circumcircle and circumcenter are visually indistinguishable from the boundary of the maximal disk and the medial axis bifurcation point. On the right, we increase the sampling density. The triangle vertices approach the contact points of the disk, and the triangle's circumcenter moves even closer to the bifurcation point.

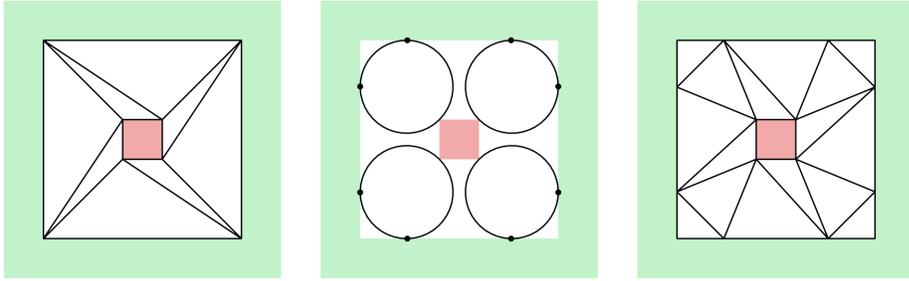


FIGURE 6.7. On the left, we show set R and G in red and green, respectively, with a constrained Delaunay triangulation of the space between R and G in black. In the middle, we outline the maximal disks centered on the bifurcation points of the medial axis of $(R \cup G)^c$ in black, and we indicate the points where the disks touch the boundary of G with black dots. On the right, we show the result of inserting those points into the edges of G as vertices and computing a constrained Delaunay triangulation.

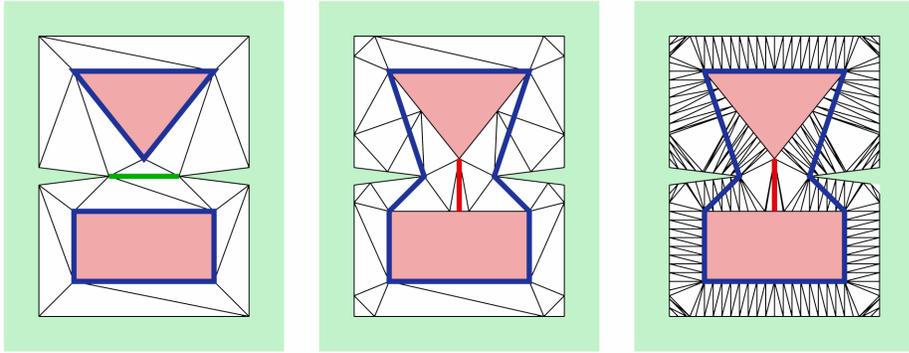


FIGURE 6.8. On the left, we show a constrained Delaunay triangulation of vertices located at the corners of R and G . We show the cover obtained from this triangulation in blue. In the middle, we add additional vertices at the contact points of the maximal disks centered on the bifurcation points of the medial axis of $(R \cup G)^c$. The cover obtained from the augmented vertex set is different from the cover obtained from the cover on the left. On the right, we densely sample the edges bounding R and G while retaining the vertices at bifurcation disk contact points. Although the vertices we add by sampling change the triangulation, they affect neither the triangles inscribed in bifurcation disks nor the computed cover.

Rather than compute an r -tightening, we compute an approximate medial tightening. The medial r -tightening of a set is the medial cover of its r -opening relative to the r -opening of its complement. Due to the algebraic complexity of tightening with

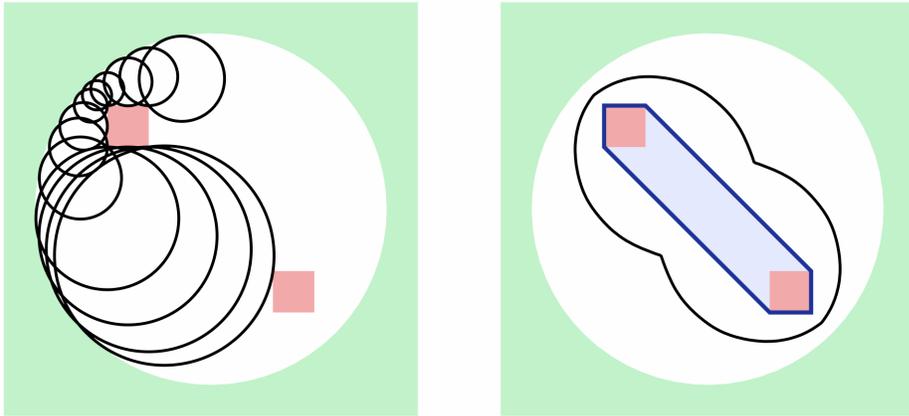


FIGURE 6.9. If we have an algorithm for computing the medial cover, we can compute the convex hull of a nonempty set R by taking G to be the complement of any closed disk containing R . Then no maximal disk in the gap between R and G touches G in more than one point. We illustrate this on the left, shading R red and G green and outlining several maximal disks in black. The set of points equidistant from R and G is a simple closed curve separating R and G , drawn in black on the right. The boundaries of the convex hull of R and the medial cover of R relative to G are both isotopic to this curve when we hold R and G fixed. As a result, the two sets are identical. We show them in blue on the right.

disks (Chapter 16,) we substitute a regular n -sided polygon P inscribed in a disk of radius r for the disk used in morphological r -opening. The cost of computing the polygonal opening with P is linear in n (Figure 6.10,) while the error between P and its circumscribing disk is proportional to $\frac{1}{n}$. However, we note that the Hausdorff distance between the r -opening of a set and its opening with P may be arbitrarily large (Figure 6.11.) Although an approximate medial tightening does not have bounded curvature, it does have bounds on the maximum angle change at a vertex and over a boundary interval (Figure 6.12.)



FIGURE 6.10. On the left, we show a square in gray. In the middle, we zoom in on a corner of the square. Overlaid on the corner in black, we show the opening of the square by a regular polygon with 10 sides. On the right, we show the opening of the square by a regular polygon with 20 sides. When we double the number of sides in the polygon, the number of edges at a corner of the square also doubles. While the number of sides in the square remains constant, the number of sides in the opening is linear in the number of sides in the polygon used for opening.

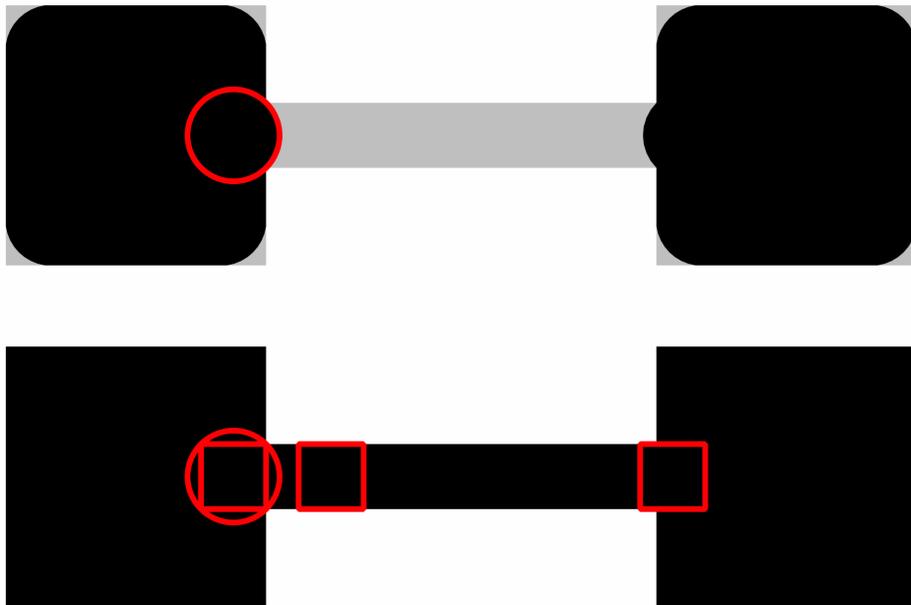


FIGURE 6.11. Above, we show a set S in gray with the opening of S using a disk of radius r overlaid in black. We outline the disk in red. On the right, we show the opening of S using a square inscribed in the disk. Unlike the opening of S with the disk, the opening of S with the inscribed square is identical to S . We can make the Hausdorff distance between the opening of S with a disk and the opening of S with a square inscribed in the disk arbitrarily large by lengthening the neck connecting the left and right halves of S .

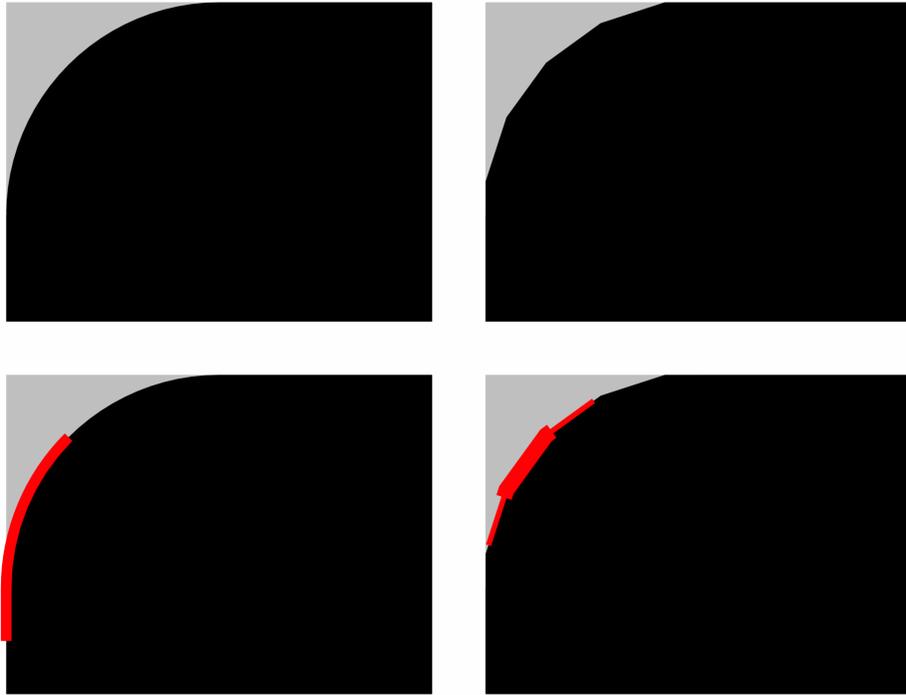


FIGURE 6.12. In the upper left, we illustrate how r -tightening replaces a corner of a square with a circular arc of radius r . In the upper right, we show how polygonal tightening replaces the corner with a piece of a regular polygon. A total angle change of $\frac{\pi}{2}$ occurs between the endpoints of the circular arc and the endpoints of the polygonal arc. Along the circular arc, the rate of change in the normal's direction is $\frac{1}{r}$, while along the polygonal arc the normal angle changes by $\frac{\pi}{10}$ at each of the arc's 5 vertices. In the lower left, we highlight an interval over the boundary of the r -tightening in red. The change in angle over an arclength interval of an r -tightening is no greater than $\frac{1}{r}$ times the interval's length. In the lower right, we indicate two overlapping intervals from the boundary of the polygonal r -tightening with two overlapping red lines, one thick and one thin. An angle change of $\frac{\pi}{5}$ occurs over both intervals because both intervals contain two arc vertices. The change in angle over an interval from the boundary of a polygonal r -tightening is no greater than r times the interval's length plus an error inversely proportional to n .

CHAPTER 7

APPLICATIONS

7.1. Convergent boundary estimation

The relative convex hull provides a convergent estimate of the measure of the boundary of a set given its rasterization on a regular lattice [58]. We conjecture that the tight hull does as well. For instance, suppose we are given a bounded, planar set S that is morphologically r -regular: the boundary of S is a manifold, and its interior and exterior can both be expressed as unions of balls radius r (Figure 7.1; Definition 8.27.) We sample S on a square grid, and we color all grid points in S red and all grid points outside of S green. We color red every edge connecting two red points and

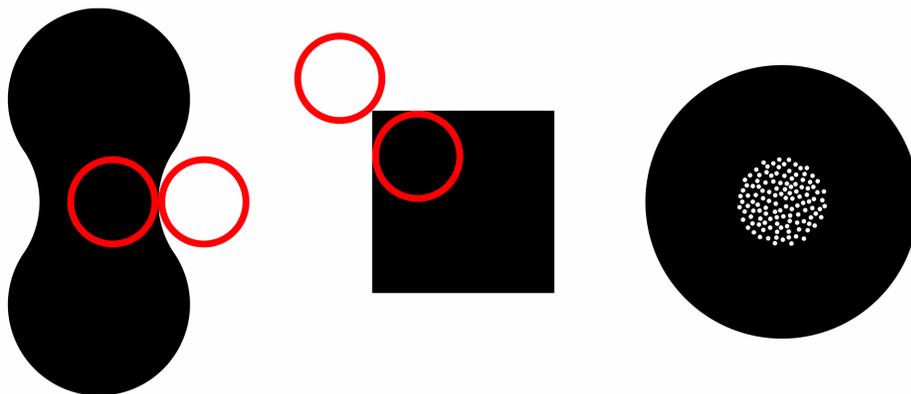


FIGURE 7.1. On the left, we show a morphologically r -regular set. It has a manifold boundary and its interior and complement can be expressed as unions of balls of radius r . The interior of the square in the middle cannot be expressed as a union of balls of radius r , so it is not r -regular. The set S on the right is the set difference between a disk of radius $3r$ and the intersection of a concentric disk of radius r with the set of points with rational coordinates. Because the boundary of S includes the disk containing the cloud of points with rational coordinates, it is not a manifold curve and S is not morphologically r -regular.

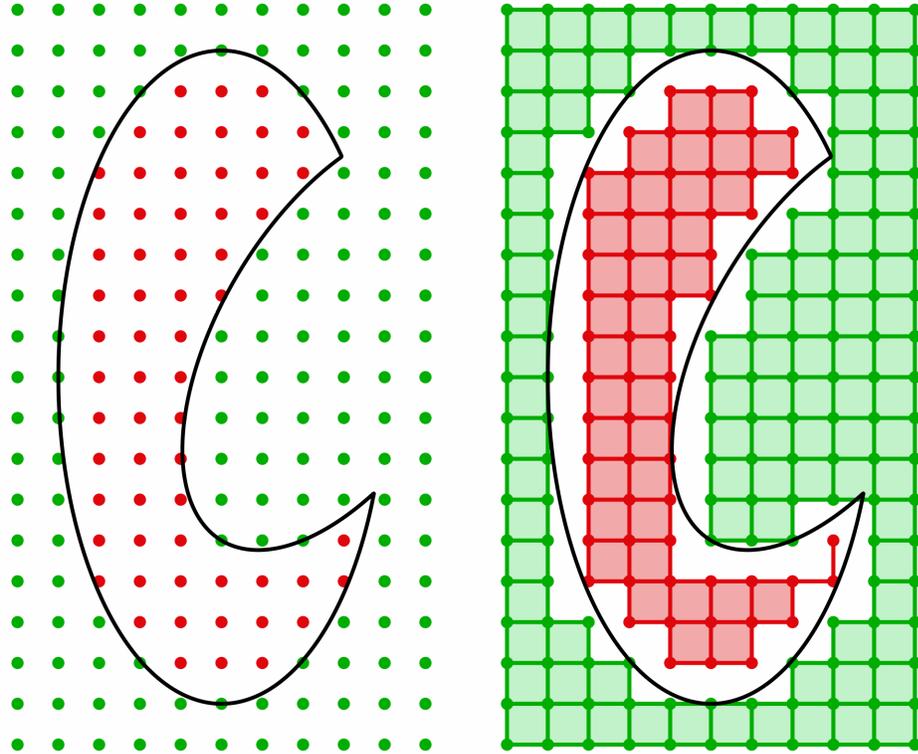


FIGURE 7.2. On the left, we outline a set S in black. We show a rectangular grid of samples, with samples inside of a set S colored red and samples outside of S colored green. On the right, we shade red the edges and squares with red vertices, while we shade green the edges and squares with green vertices.

every square bounded by four red edges. We likewise color green the edges connecting green vertices and the squares bounded by green edges (Figure 7.2.)

Sloboda and Zatko prove that when the spacing between samples goes to zero, the measure of the boundary of the convex hull of the red set R relative to the green set G converges to the measure of the boundary of S (Figure 7.3; [58].) In Chapter 18, we show that in this situation the tight hull of R relative to G is identical to the convex hull of R relative to G . As a result, the boundary estimates from the tight hull and relative convex hull agree, implying the tight hull provides a convergent boundary measure estimate in two dimensions.

While Peano and Jordan [30, 45] obtained convergent estimates of the d -dimensional measure of a d -dimensional set in the 19th century (Figure 7.4,) convergent $(d - 1)$ -

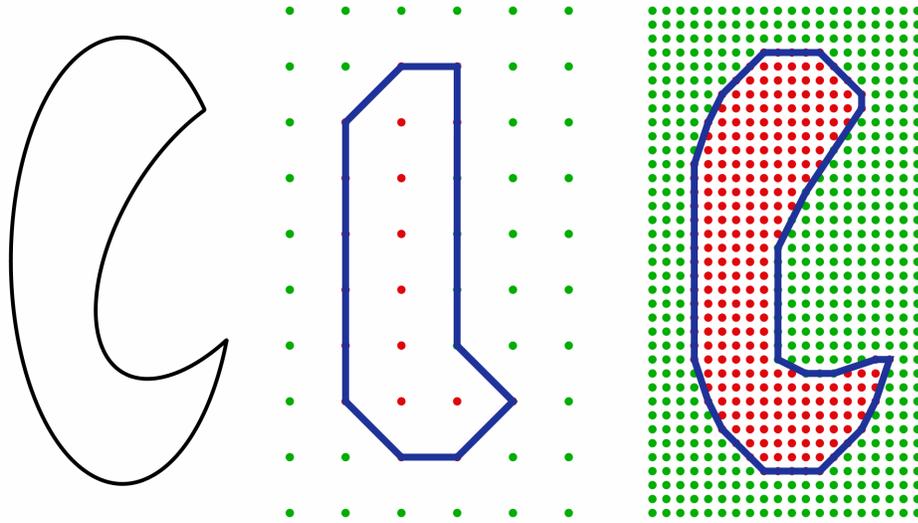


FIGURE 7.3. On the left, we outline a set S with boundary length of 91.0. In the middle, we show the relative convex hull's reconstruction of the boundary of S from a coarse sampling in blue. The length of the reconstructed boundary is 70.6. On the right, we show the relative convex hull's boundary reconstruction from a fine sampling, which has a length of 84.2. The length of the reconstructed boundary converges to the length of the boundary of S as the spacing between grid samples goes to zero.

dimensional boundary measure estimation has proved more difficult. In part this is because we can expect surfaces extracted from volume data by assembling patches constructed from samples in local windows to have nonconvergent areas. Consider the two-dimensional midpoint reconstruction, which we define as the result of placing a vertex at the midpoint of every edge and then connecting vertices whose edges bound the same square. If the boundary of the set to be rasterized includes a line that is not aligned with the rasterization grid, the midpoint reconstruction exhibits staircasing (Figure 7.5.)

Although we can reduce the resulting length error by constructing boundary segments using data drawn from larger windows, there is no finite window size such that the error vanishes as the intersample spacing vanishes (Chapter 18.) Because the midpoint reconstruction overestimates the length of line segments that are not aligned with the lattice axes, boundary measure estimates obtained from it contradict

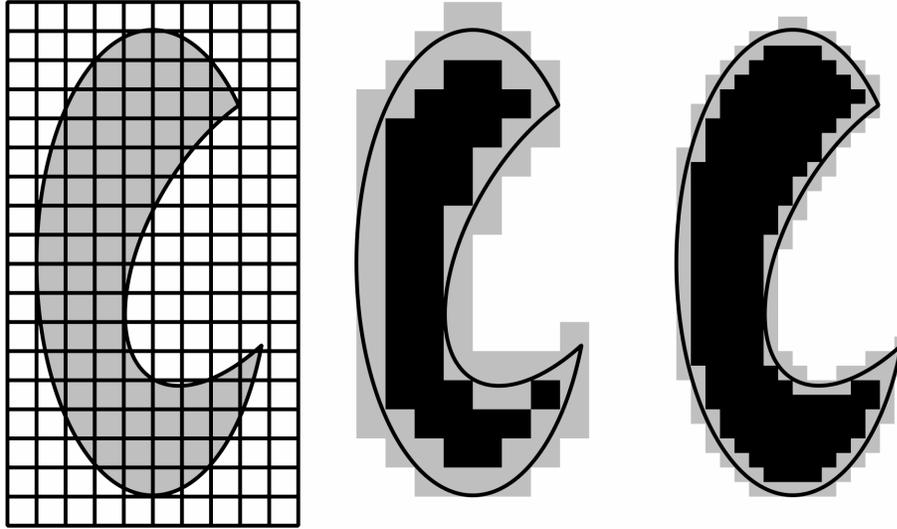


FIGURE 7.4. Peano and Jordan define the area of a planar set S by overlaying a square grid on it, as illustrated on the left. Taking the area of a square to be the product of its edge lengths, they define the outer measure of S as the sum of the areas of the squares intersected by S , while the inner measure is the sum of the areas of the squares contained in S . In the middle, we shade the squares intersected by S gray with the squares contained in S overlaid in black. If the outer measure of S approaches the inner measure of S as the edge length of the grid's squares goes to zero, then S is measurable, and its measure (or area in this case, because S is two-dimensional) is their common limit. We illustrate the computation of inner and outer measures using short grid edges on the right.

fundamental theorems of geometry, such as the proposition that a circle is the set of minimum perimeter out of all sets with a given area [34].

The tight hull and the relative convex hull both cut across lattice cells, rather than forming the staircase patterns seen in the midpoint reconstruction. Staircasing causes slack to increase as intersample spacing decreases. Because the tight hull minimizes slack, we instead expect its slack to be convergent. If nonconvergent boundary measure estimators have excess slack and the tight hull's slack is convergent, then the tight hull is a convergent boundary measure estimator.

Convergent boundary estimators are valuable in the numerical solution of definite integration problems [34]. For example, the surface area of an ellipsoid or a Boolean

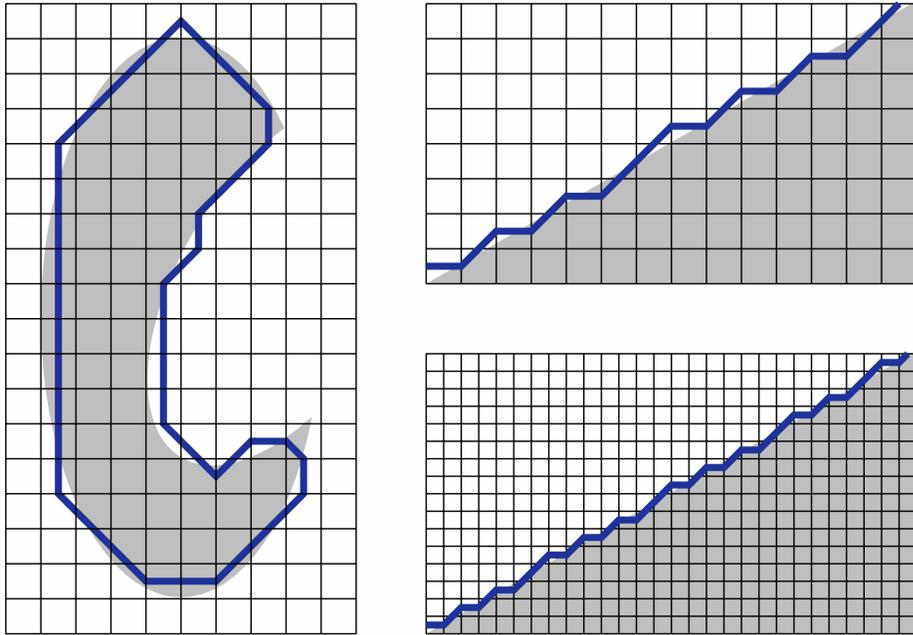


FIGURE 7.5. On the left, we show a midpoint reconstruction of the boundary of the set S from Figure 7.3 in blue. Compared to the boundary reconstruction obtained from the relative convex hull, the midpoint reconstruction is more jagged, which increases its length. To highlight this behavior, we consider the rasterization of the edge of a set lying at an angle of $\frac{\pi}{6}$ to the horizontal rasterization axis. If the edge of the gray set samples on the upper right has unit length, then the midpoint reconstruction shown has a length of 1.06. Halving the intersample spacing as shown on the lower right does not reduce the midpoint reconstruction's length, illustrating that it is a nonconvergent boundary length estimator.

combination of ellipsoids cannot generally be expressed in terms of elementary functions, but rasterizing ellipsoids is straightforward. Once we have a rasterization, we can construct a boundary estimate from the volume samples whose measure converges to the rasterized set's boundary measure as the intersample spacing goes to zero. Assuming tight hulls constructed from polygonal data are polyhedral, we express the hull boundary as a triangle mesh and compute its area by summing the areas of its triangles.

The result of that summation is a single number. While this may be all that many problems require, the reconstructed surface itself may be valuable in some applications. For instance, suppose our input is a segmented medical image, where each

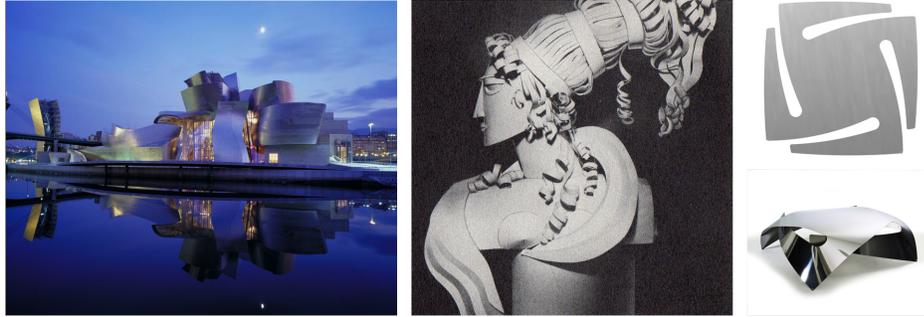


FIGURE 7.6. On the left, we show an image of the facade of the Guggenheim at Bilbao [?]. In the middle, we show a paper sculpture used in a perfume store display [25], and on the right we show a coffee table folded from a laser-cut sheet of stainless steel [31].

voxel is labeled as inside or outside of a given anatomical structure. The shading of a surface we extract from the image is determined by its normals, so the surface's normal field ideally converges to the normal field of the segmented structure. Given that the normal fields of tight hulls exhibit minimal variation, we can reasonably ask whether their normal fields are convergent. If tight hull normal fields are convergent, tight hulls provide a way to visualize segmented medical images with a formal guarantee that the visualization faithfully depicts the data.

7.2. Shape design

Tight hulls may prove useful in shape design. It appears that the portion of a tight hull disjoint from the boundaries of the sets defining it consists of developable patches. Developable patches have properties that facilitate their representation and manufacture, and a body of literature discusses their application to modeling in fields such as architecture and industrial design (Figure 7.6; [24, 32, 48, 50].) While several papers address the problem of computing a developable patch interpolating a given curve (Figure 7.7; [24, 50],) tight hulls may be the first work on developable surfaces with volumetric constraints.

Provided an efficient algorithm for three-dimensional tight hull construction, we can envision interactively modeling with tight hulls by manipulating the sets they

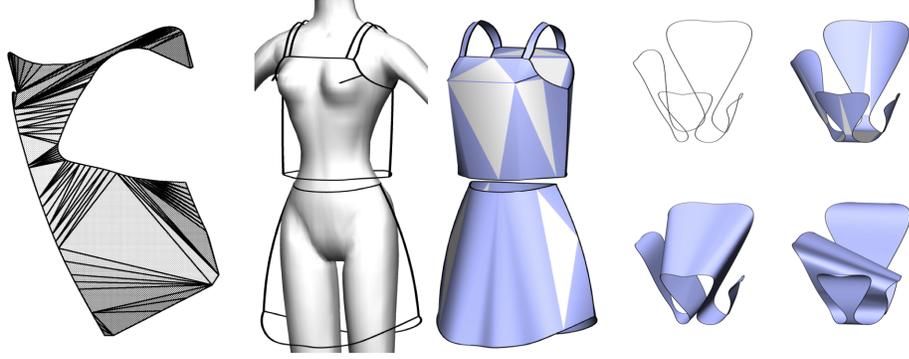


FIGURE 7.7. On the left, we show a triangulation of a boundary curve [24] that estimates the geometry assumed by a sheet of metal pressed by a machine. The center image [50] illustrates garment design by interpolating seams with developable patches. The right images [50] show how multiple patches can interpolate the same boundary curve, providing us with motivation to determine which patch is “best.”

include and exclude. For instance, we could produce an executable illustration of the thought experiment described in Subsection 3.1, where we press a finger into the boundary of a hollow cylinder’s convex hull. Presently, however, such a simulation is feasible only if all the sets involved share rotational symmetry about the same axis. At an intermediate level of complexity, we might interactively manipulate the tight hulls of height fields. This could provide a means for image editing and representation akin to diffusion curves [44], but with linear gradients along the rulings of the developable surfaces we construct.

7.3. Blending

We anticipate two significant applications of tightening. The first is to solid blending. In two dimensions, tightening a set limits the magnitude of its curvature, making it globally smooth. By contrast, it is impossible to make a set smooth with morphological opening or closing, or with any composition of the two (Chapter 14.) At present, we know little about the properties of tightening in three and higher dimensions, but tightening apparently guarantees the existence of a bounded sectional curvature at every boundary point. Even in the absence of a guaranteed curvature bound, tightening may produce useful blends in three dimensions, just as opening

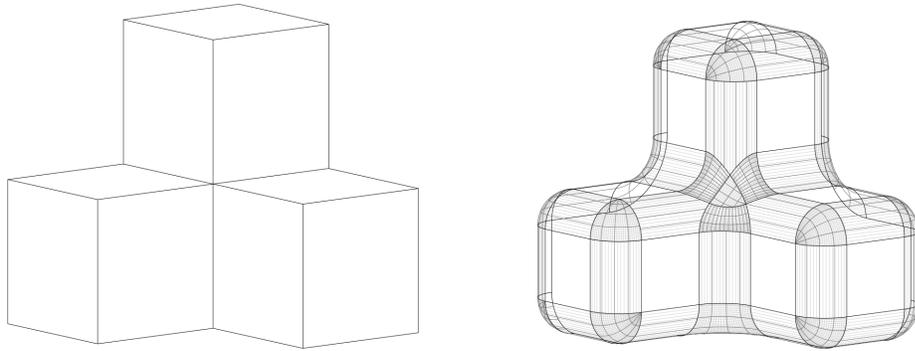


FIGURE 7.8. On the left, we show a wireframe image of a solid S with a singular saddle point. On the right, we show a schematic illustration of the tightening of S with a small radius, which has a unique normal at every point.

and closing produce useful blends despite the fact that do not guarantee smoothness. Compared to a composition of opening and closing, tightening has the advantage that it is symmetric with respect to set complement, so it has no bias toward either enlarging or reducing a set.

7.4. Normal field simplification

Tightening's second application is to normal field simplification. Tightening has potential for use as a theoretical tool. Although we define normal fields to sets with singularities in Section 10.3, the outward-pointing normal at a singularity is not unique. A small-radius tightening of a set, however, is a close approximation to the set with a unique normal at every point on its boundary (Figure 7.8.) As we increase the tightening radius, we obtain scale-dependent versions of the set's normal field, which we can visualize by the coverage of the image of the normals on a sphere (Figure 7.9.)

7.5. Further applications

Some applications of tightening show promise but remain speculative. An example is the use of tightening for the topological repair of polygonal models. In two dimensions, suppose we are given an unorganized collection of edges and vertices. We

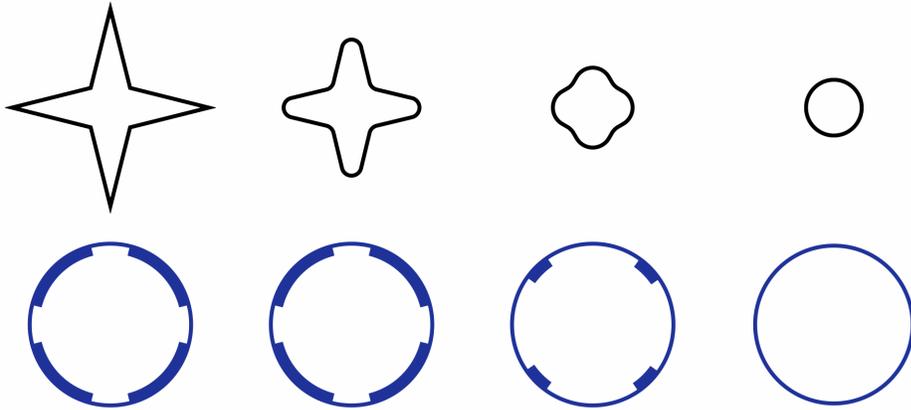


FIGURE 7.9. On the top row, we show the progressive tightening of a four-pointed star. Beneath each star we show the star's normal field coverage. The untightened and slightly tightened stars have identical normal fields, but the normal field variation decreases as the tightening radius further increases. Before disappearing entirely, the tightened star is a disk, and its normal field uses every normal direction exactly once.

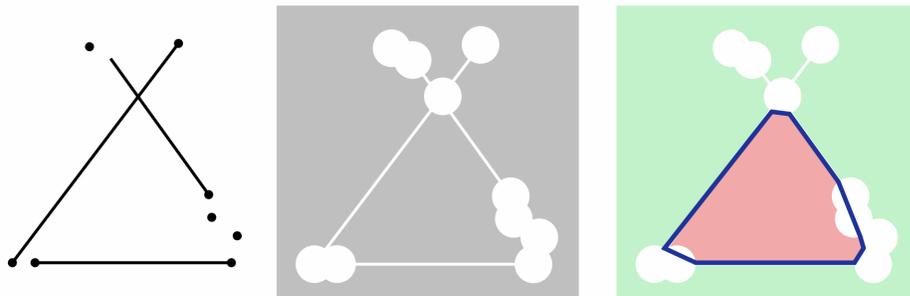


FIGURE 7.10. On the left, we show a nonmanifold soup of edges and vertices. In the middle, we shade gray the set X obtained by subtracting the union of the input with the dilation of its nonmanifold points from the plane. On the right, we designate the bounded component of X as belonging to the red set R and the unbounded component as belonging to the green set G . Our repaired model is the tight hull of R relative to G , outlined in blue.

intersect the edges in the data, then identify nonmanifold vertices. We dilate the nonmanifold vertices and subtract them as well as the input set from the plane. We take the tight hull of the bounded components relative to the unbounded component to produce a repaired model (Figure 7.10.) Due to the properties of tight hulls, the bounding loops of the result are manifold polygons.

This plan raises questions concerning its correctness and the feasibility of its extension to three and higher dimensions. However, it illustrates a principle employed in the definition of a tightening: a tight hull extracts a set with a maximally convex boundary from a tolerance zone. We can apply tight hulls to a variety of problems by varying the tolerance zone.

PART 2

BACKGROUND

CHAPTER 8

CONVEX STRUCTURES AND CLOSURE STRUCTURES

8.1. Overview

We begin by providing a concise notation for specifying sets of open balls. We then use this notation to define the interior, closure, and boundary of a subset of Euclidean space. Exchanging balls of arbitrarily small radius for balls of a fixed, finite radius, we obtain morphological opening, morphological closing, and the mortar. We express morphological opening and closing as compositions of dilation and erosion, which under our definitions specialize Minkowski sum and difference. We then define convex hulls, relative convex hulls, geodesic convex hulls, and affine hulls for subsets of Euclidean space.

We define closure structures in terms of closed sets, whose distinguishing property is that the intersection of closed sets is closed. We then define the closure, interior, and boundary operators associated with a closure structure, and we additionally define terms for describing the operators' algebraic properties. We define convex structures in terms of convex sets, which are stable under nested union in addition to being closed under intersection. We distinguish convex hull operators that satisfy the exchange axiom from those that satisfy the antiexchange axiom [63]: the former behave like the Euclidean affine hull while the latter behave like the Euclidean convex hull.

We close by defining morphological and topological regularity, as well as symmetry with respect to set complement. We characterize topological regularity as a symmetric form of the open-regularity and closed-regularity used in solid modeling, and we characterize morphological regularity as a symmetric form of being morphological open or closed.

8.2. Ball notation

Summary. We introduce a notation for specifying sets of balls that lie in a given set, whose centers lie in a given set, or whose radii satisfy a given binary relation.

Definition 8.1. Given a set X , the set $R \subseteq X \times X$ is a *binary relation* on X .

Notation 1. For a binary relation R on X , it is common to use an infix operator to indicate that an ordered pair is an element of the relation. For instance, if we associate the symbol \leq with R , then for $a, b \in X$, $a \leq b \Leftrightarrow (a, b) \in R$.

Notation 2. Let $B_r(p)$ denote the d -dimensional open ball of radius $r \in \mathbb{R}^>$ centered at point $p \in \mathbb{R}^d$:

$$B_r(p) := \{q \in \mathbb{R}^d \mid \|q - p\| < r\}$$

For an open ball b , let $\rho(b)$ denote the radius of b and let $\pi(b)$ denote its center.

Let \mathbb{B} denote the set of all d -dimensional open balls:

$$\mathbb{B} := \{B_r(p) \mid (r \in \mathbb{R}^>) \wedge (p \in \mathbb{R}^d)\}$$

Let \mathbb{B}^S denote the set of all d -dimensional balls contained in $S \subseteq \mathbb{R}^d$:

$$\mathbb{B}^S := \{b \in \mathbb{B} \mid b \subseteq S\}$$

Let \sim be a binary relation on real numbers. Then for $r \in \mathbb{R}^>$, let $\mathbb{B}_{\sim r}$ be denote the set of all open balls $b \in \mathbb{B}$ such that $\rho(b) \sim r$:

$$\mathbb{B}_{\sim r} := \{b \in \mathbb{B} \mid \rho(b) \sim r\}$$

Let $\mathbb{B}(S)$ denote the set of open balls with centers lying in $S \subseteq \mathbb{R}^d$:

$$\mathbb{B}(S) := \{b \in \mathbb{B} \mid \pi(b) \in S\}$$

For $p \in \mathbb{R}^d$, let $\mathbb{B}(p) := \mathbb{B}(\{p\})$.

Let $\mathbb{B}_r := \mathbb{B}_{=r}$, and let $\mathbb{B}_{\sim r}^X(Y) := \mathbb{B}^X \cap \mathbb{B}_{\sim r} \cap \mathbb{B}(Y)$.

Example 8.1. Using this notation, $\mathbb{B}_{>r}$ is the set of open balls with radii greater than r , and $\mathbb{B}_{>r}^S$ is the set of open balls contained in S with radii greater than r . Similarly, \mathbb{B}_r^S is the set of all open balls of radius r contained in S . The union of the balls in $\mathbb{B}_r(\partial S)$ is the set of points within distance r of the boundary of S :

$$\bigcup \mathbb{B}_r(\partial S) = \{p \in \mathbb{R}^d \mid d(p, \partial S) < r\}$$

8.3. Interior, closure, opening, and closing; duality

Summary. We define topological closure, interior, and boundary operators for Euclidean space using our ball notation. Restricting ourselves to balls of radius $r \in \mathbb{R}^>$, we provide analogous definitions for morphological opening, morphological closure, and the mortar.

8.3.1. Topological operators.

Remark 8.1. We provide definitions for the standard topological operators for induced by the Euclidean metric. Note that a $(d - 1)$ -dimensional set in \mathbb{R}^d is contained in its own boundary. For example, the boundary of a line segment in \mathbb{R}^2 is the whole line segment, not its endpoints.

Definition 8.2. For $S \subseteq \mathbb{R}^d$, the *interior* of S , denoted S° , is defined as

$$S^\circ := \bigcup \mathbb{B}^S$$

The *closure* of S , denoted S^- , is defined as

$$S^- := S^{\text{coc}}$$

The *boundary* of S , denoted ∂S , is defined as

$$\partial S := S^- \cap S^{c-}$$

Remark 8.2. The interior of a set is the union of all open balls contained within it. The closure of a set is the complement of the interior of its complement, and the boundary of a set is the intersection of the set's closure with the closure of its complement.

8.3.2. Morphological operators.

Summary 1. We define morphological opening, closing, and mortar. Informally, the r -opening of a set is the set of points swept by an open ball of radius r inside the set, the r -closing of a set is the complement of the set of points swept by a ball of radius r rolling outside the set, and the r -mortar of a set is points that cannot be reached by a ball disjoint from the set's boundary. Section 4.1 contains an illustrated explanation.

Definition 8.3. For $S \subseteq \mathbb{R}^d$ and $r \in \mathbb{R}^>$, the *morphological opening* of S with radius r , denoted $S \circ_r$, is defined as:

$$S \circ_r := \bigcup \mathbb{B}_r^S$$

The *morphological closure* of S with radius r , denoted $S \bullet_r$, is defined as

$$S \bullet_r := S^c \circ_r^c$$

The r -*mortar* of S , denoted $\mathcal{M}_r S$, is defined as

$$\mathcal{M}_r S := S \bullet_r \cap S^c \bullet_r$$

Remark 8.3. We define the morphological opening of a set with radius r as the union of open balls of radius r contained within it. The morphological closure of a set with radius r is the complement of the opening of its complement with radius r ,

and a set's mortar is the intersection of a set's morphological closure with radius r with the closure of its complement.

8.3.3. Duality.

Remark 8.4. The relationship of the interior operator to the closure operators is the same as the relationship of morphological opening to morphological closing: $S^- := S^{\text{coc}}$ and $S \bullet_r := S^c \circ_r^c$. We formalize this relationship by defining duality.

Definition 8.4. Operators $\Psi : \wp(X) \rightarrow \wp(X)$ and $\Upsilon : \wp(X) \rightarrow \wp(X)$ are *dual* if and only $\Psi(Y^c) = \Upsilon(Y)^c$.

Remark 8.5. Equivalently, Ψ and Υ are dual if and only if $\Psi(X) = \Upsilon(X^c)^c$. Topological interior and closure are dual, as are morphological opening and closing.

8.4. Dilation, erosion, and Minkowski sum and difference

Summary. We formulate definitions of dilation and erosion using open balls, and we express morphological opening and closing as compositions of erosion and dilation. We define Minkowski sums and differences, which generalize dilation and erosion, and we use Minkowski operators to define generalized forms of opening and closing. In particular, we define polygonal forms of opening and closing where a regular polygon replaces the open disk of radius r used in morphological r -opening and r -closing.

8.4.1. Dilation, erosion, and their compositions.

Remark 8.6. Dilating a set by $r \in \mathbb{R}^>$ adds to it all points within a distance r of it, increasing its size. Similarly, erosion decreases a set's size by removing all points within a distance of its complement. As we explain in Remark 8.7, if we dilate a set by r and then erode it by r , we obtain the set's r -closure. Similarly, if we erode by r and then dilate by r , we obtain r -opening.

Definition 8.5. For $S \subseteq \mathbb{R}^d$ and $r \in \mathbb{R}^>$, the *dilation* of S by r , denoted $S \uparrow_r$, is defined as

$$S \uparrow_r := \bigcup \mathbb{B}_r(S)$$

The *erosion* of S by r , denoted $S \downarrow_r$, is the dual of dilation:

$$S \downarrow_r := S^c \uparrow_r^c$$

The *r-offset zone* of S , denoted $\mathcal{Z}_r S$, is defined as

$$\mathcal{Z}_r S := S \uparrow_r \cap S^c \uparrow_r$$

Remark 8.7. We can express morphological opening and closing as compositions of erosion and dilation:

$$S \circ_r = S \downarrow_r \uparrow_r$$

$$S \bullet_r = S \uparrow_r \downarrow_r$$

The erosion of $S \downarrow_r$ of $S \subseteq \mathbb{R}^d$ by $r \in \mathbb{R}^>$ is the set of centers of all open balls of radius r lying in S . When we dilate $S \downarrow_r$ to obtain $S \downarrow_r \uparrow_r$, we dilate the centers of the open balls of radius r lying in S . This yields the union of the open balls of radius r in S , which is precisely the opening $S \circ_r$. Similar reasoning demonstrates $S \bullet_r = S \uparrow_r \downarrow_r$.

Remark 8.8. The *r*-offset zone $\mathcal{Z}_r S$ is the set of all points within a distance r of ∂S :

$$\mathcal{Z}_r S = (\partial S) \uparrow_r = \bigcup \mathbb{B}_r(\partial S) = \{p \in \mathbb{R}^d \mid d(p, \partial S) < r\}$$

If the Hausdorff distance between two sets $A, B \subseteq \mathbb{R}^d$ is less than r , then $\partial A \subseteq \mathcal{Z}_r B$ and $\partial B \subseteq \mathcal{Z}_r A$. However, the Hausdorff distance between A and B is less than r if and only if $A \subseteq B \uparrow_r$ and $B \subseteq A \uparrow_r$. The difference between the conditions is that if A and B are within a Hausdorff distance of r , then A^c and B may not be within a Hausdorff distance of r , despite the fact that both

$$(\partial A \subseteq \mathcal{Z}_r B) \wedge (\partial B \subseteq \mathcal{Z}_r A)$$

For instance, the Hausdorff distance between a bounded set A and its complement is infinite, but for all $r \in \mathbb{R}^>$, we have $\partial A \subseteq \mathcal{Z}_r A^c$ and $\partial A^c \subseteq \mathcal{Z}_r A$.

8.4.2. Minkowski sum and difference.

Summary 2. We provide definitions of the Minkowski sum and difference, which generalize our definitions of dilation and erosion. We can interpret dilation as a “logical,” “binary,” or “set-theoretic” convolution with a ball-shaped kernel. Minkowski sums similarly resemble convolution, but do not necessarily use balls as kernels.

Notation 3. For $S \subseteq \mathbb{R}^d$ with origin $o \in \mathbb{R}^d$, let \tilde{S} denote the reflection of S about the origin:

$$\tilde{S} := \{o + (o - p) \mid p \in S\}$$

Definition 8.6. For $A, B \subseteq \mathbb{R}^d$, suppose we are given $o \in \mathbb{R}^d$ as the origin for B . Then the *Minkowski sum* of A and B , denoted $A \oplus B$, is defined as

$$A \oplus B := \bigcup_{p \in B} A + (p - o)$$

The *Minkowski difference* of A and B , denoted $A \ominus B$, is defined as

$$A \ominus B := \bigcap_{p \in \tilde{B}} A + (p - o)$$

Remark 8.9. If we select a point as the origin for A , then $A \oplus B = B \oplus A$. By convention, however, the set B in a Minkowski sum or difference of the form $A \oplus B$ is referred to as the *structuring element*. The set A is the image or data to be processed, while B is analogous to a convolution kernel.

8.4.3. Dual opening and closing from adjoint erosion and dilation.

Summary 3. We present the construction of dual opening and closing operators from adjoint dilation and erosion operators defined in terms of Minkowski operators [28, 55].

Definition 8.7. Operators $\delta : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\varepsilon : \mathbb{R}^d \rightarrow \mathbb{R}^d$ are *adjoint* if and only if for all $X, Y \subseteq \mathbb{R}^d$,

$$X \subseteq \varepsilon(Y) \Leftrightarrow \delta(X) \subseteq Y$$

Definition 8.8. For $B \subseteq \mathbb{R}^d$, suppose we are given $o \in \mathbb{R}^d$ as the origin for B . Then the *B-dilation* $\delta_B : \wp(\mathbb{R}^d) \rightarrow \wp(\mathbb{R}^d)$ is defined as $\delta_B(S) := S \oplus B$.

The *B-erosion* $\varepsilon_B : \wp(\mathbb{R}^d) \rightarrow \wp(\mathbb{R}^d)$ is defined as $\varepsilon_B(S) := S \ominus B$.

Remark 8.10. The operators δ_B and ε_B are adjoint, while δ_B and $\varepsilon_{\tilde{B}}$ are dual. In particular, if B is symmetric about the origin then δ_B and ε_B are dual. For example, the dilation and erosion operators \uparrow_r and \downarrow_r are both adjoint and dual, because the open ball of radius r is symmetric about the origin.

The compositions $\delta_B \varepsilon_{\tilde{B}}$ and $\varepsilon_{\tilde{B}} \delta_B$ are also dual:

$$\begin{aligned} \delta_B(\varepsilon_{\tilde{B}}(S^c))^c &= \left((S^c \ominus \tilde{B}) \oplus B \right)^c = \left((S^c \ominus \tilde{B})^c \ominus \tilde{B} \right) \\ &= \left((S \oplus B) \ominus \tilde{B} \right) = \varepsilon_{\tilde{B}}(\delta_B(S)) \end{aligned}$$

If B is symmetric about the origin, $\delta_B \varepsilon_{\tilde{B}} = \delta_B \varepsilon_B$ and $\varepsilon_{\tilde{B}} \delta_B = \varepsilon_B \delta_B$ are dual. As explained in Section 8.6, $\downarrow_r \uparrow_r$ and $\varepsilon_{\tilde{B}} \delta_B$ are both closure operators that define closure structures. Just as $\downarrow_r \uparrow_r$ is the interior operator for the closure structure defined by $\uparrow_r \downarrow_r$, the composition $\delta_B \varepsilon_{\tilde{B}}$ is the interior operator for the closure structure defined by $\varepsilon_{\tilde{B}} \delta_B$.

Theorem 8.1. (Serra [55]) For $S \subseteq \mathbb{R}^d$, the composition $\delta_B (\varepsilon_{\tilde{B}} (S))$ is equal to the union of all translated copies of B that lie within S :

$$\delta_B (\varepsilon_B (S)) := \bigcup_{\{p \in \mathbb{R}^d \mid B+p \subseteq S\}} B + p$$

PROOF. The set

$$\varepsilon_{\tilde{B}} (S) = \bigcap_{p \in \tilde{B}} S + (p - o) = \bigcap_{p \in B} S + (o - p)$$

If $q \in \varepsilon_{\tilde{B}} (S)$, then $q + (p - o) \in S$, while if $q + (p - o) \in S$, then $q \in \varepsilon_{\tilde{B}}$. Then

$$\begin{aligned} \varepsilon_{\tilde{B}} (S) &= \{q \in \mathbb{R}^d \mid \forall p \in B : q + (p - o) \in S\} \\ &= \{q \in \mathbb{R}^d \mid B + q \subseteq S\} \end{aligned}$$

Because

$$\delta_B (S) = S \oplus B = B \oplus S$$

$$\delta_B (\varepsilon_{\tilde{B}} (S)) = \bigcup_{p \in B} \varepsilon_{\tilde{B}} (S) + (p - o) = \bigcup_{q \in \varepsilon_{\tilde{B}} (S)} B + q = \bigcup_{\{q \in \mathbb{R}^d \mid B+q \subseteq S\}} B + q$$

We conclude that the composition $\delta_B (\varepsilon_{\tilde{B}} (S))$ is the union of all translations of B contained in S . □

8.4.4. Polygonal opening and closing.

Remark 8.11. If $B \subseteq \mathbb{R}^2$ is a regular polygon symmetric about the origin inscribed in a circle of radius r , then δ_B is a polygonal form of dilation by r , while ε_B is a polygonal form of erosion by r .

Definition 8.9. For $S \subseteq \mathbb{R}^2$, let $P \subseteq \mathbb{R}^2$ be a regular polygon with an even number k of sides inscribed in an open disk of radius r . We specify an orientation for P by assigning a coordinate system to the plane such that P is aligned with that coordinate system so that its origin coincides with P 's center and the y -axis bisects one of P 's edges. Then the (k, r) -opening of S , denoted \circ_r^k , is defined as

$$S \circ_r^k := \bigcup_{\{q \in \mathbb{R}^2 \mid P+q \subseteq S\}} P + q$$

The (k, r) -closing of S , denoted $S \bullet_r^k$, is defined the dual of the (k, r) -opening of S :

$$S \bullet_r^k := S^c \circ_r^{k,c}$$

The (k, r) -mortar of S , denoted $\mathcal{M}_r^k S$, is defined as

$$\mathcal{M}_r^k S := S \bullet_r^k \cap S^c \bullet_r^k$$

Remark 8.12. Our definition of the (k, r) -opening is unique because we specify an orientation for P . Each distinct rotation of P about the origin yields a different closure, closing, and mortar operator.

8.5. Convex hulls, relative convex hulls, and affine hulls

Summary. We provide definitions of the Euclidean convex hull, the relative convex hull, the geodesic hull, and the Euclidean affine hull.

Definition 8.10. A set $S \subseteq \mathbb{R}^d$ is *convex* if and only if for every pair of points $p, q \in \mathbb{R}^d$, the closed line segment \overline{pq} is contained in S . A set is *concave* if and only if its complement is convex.

The *Euclidean convex hull* of $S \subseteq \mathbb{R}^d$, denoted $\text{CH}(S)$, is defined as the intersection of all convex sets containing S .

Remark 8.13. Let C denote the convex subsets of \mathbb{R}^d . Then we can write the convex hull $\text{CH} : \wp(\mathbb{R}^d) \rightarrow \wp(\mathbb{R}^d)$ as

$$\text{CH}(S) = \bigcap \{X \in C \mid S \subseteq X\}$$

Different authors define convex subsets of Euclidean space and Euclidean convex hulls in different but logically equivalent ways. Definition 8.10 is typical. For an illustrated discussion of convex sets and convex hulls, see Subsection 3.2.1.

Definition 8.11. For $G \subseteq \mathbb{R}^d$, a set $S \subseteq G^c$ is *convex relative to G* if and only if for all $p, q \in S$,

$$\overline{pq} \subseteq G^c \Rightarrow \overline{pq} \subseteq S$$

The *convex hull of $S \subseteq G^c$ relative to G* , denoted $\text{CH}(S \mid G)$, is the intersection of all sets convex relative to G containing S .

Remark 8.14. This definition of the relative convex hull is based on the definition provided by Sklansky and Kibler [57], which was later adopted by Sloboda and Zatko [58]. For an illustrated discussion of the relative convex hull, see Subsection 3.2.2.

Definition 8.12. Suppose that there is a unique shortest path between every pair of points in the complement of $G \subseteq \mathbb{R}^d$. Then $S \subseteq G^c$ is *geodesically convex* relative to G if and only if for every $p, q \in S$, the shortest path connecting p and q in G^c is contained in S .

The *geodesic hull of $S \subseteq G^c$ relative to G* , denoted $\text{GH}(S \mid G)$, is the intersection of all sets geodesically convex relative to G that contain S .

Remark 8.15. We extend Toussaint's definition [62], who establishes the existence of a unique shortest path between every pair of points by assuming G^c is a simple polygon. If there is not always a unique shortest path between two points, we may

alternatively define a set as geodesically convex if and only if it contains every path of minimum length connecting any pair of its points.

Definition 8.13. A set $S \subseteq \mathbb{R}^d$ is *affine* if and only if for every pair of points $p, q \in \mathbb{R}^d$, the line \overleftrightarrow{pq} is contained in S .

The *Euclidean affine hull* of $S \subseteq \mathbb{R}^d$, denoted $\text{AH}(S)$, is the intersection of all affine sets containing S .

Remark 8.16. The affine hull of $S \subseteq \mathbb{R}^d$ is the lowest-dimensional affine subspace containing S . For instance, if $S \subseteq \mathbb{R}^3$ is a circle lying in three dimensions, the affine hull of S is the plane containing it.

8.6. Closure structures

Summary. We provide a definition of closure structures, which capture the properties common to operators like topological and morphological closure. We define properties important to operators associated with closure structures. We introduce self-duality and invariance with respect to set complement, which are operator properties related to duality.

8.6.1. Closure structures and closure operators.

Summary 4. As in van de Vel's monograph [63], we define closure structures in terms of closed sets, and we define the closure, interior, and boundary operators determined by a closure structure.

Definition 8.14. For a set X , suppose $\mathbb{C} \subseteq \wp(X)$. Then the pair (X, \mathbb{C}) is a *closure structure* if and only if

- (1) $\emptyset \in \mathbb{C}$
- (2) $X \in \mathbb{C}$
- (3) For all $\mathbb{D} \subseteq \mathbb{C}$, $\bigcap \mathbb{D} \in \mathbb{C}$.

If (X, \mathbb{C}) is a closure structure, we say the elements of \mathbb{C} are *closed*. A set $Y \subseteq X$ is *open* if and only if Y^c is closed.

Remark 8.17. The key property of a closure structure (X, \mathbb{C}) is that any intersection of closed sets is closed. Defining the set X as closed guarantees that every subset of X is contained in a closed set, while defining the empty set to be closed entails that the intersection of disjoint closed sets is closed.

Definition 8.15. For a closure structure (X, \mathbb{C}) , the *closure operator* $\Psi : \wp(X) \rightarrow \wp(X)$ is defined as

$$\Psi(Y) := \bigcap \{Z \in \mathbb{C} \mid Y \subseteq Z\}$$

The *interior operator* $\Phi : \wp(X) \rightarrow \wp(X)$ is defined as the dual of closure:

$$\Phi(Y) := \Psi(Y^c)^c$$

And the *boundary operator* $\Upsilon : \wp(X) \rightarrow \wp(X)$ is defined as

$$\Upsilon(Y) := \Psi(Y) \cap \Psi(Y^c)$$

Remark 8.18. The closure operator for a closure structure (X, \mathbb{C}) maps a set $S \subseteq X$ to the “smallest” closed set containing S , which is the intersection of all closed sets containing S . A closure operator uniquely defines a closure structure, and vice versa.

Topological and morphological closure are closure operators, while topological and morphological opening are interior operators. The topological boundary and mortar operators are boundary operators.

8.6.2. Increasing, idempotent, extensive, and antiextensive operators.

Summary 5. We define idempotent, increasing, extensive, and antiextensive operators. We also define the invariance of an operator, which is a collection of sets containing every set the operator maps to itself.

Definition 8.16. For set X , operator $\Psi : \wp(X) \rightarrow \wp(X)$ is *idempotent* if and only if for all $Y \subseteq X$,

$$\Psi(Y) = \Psi(\Psi(Y))$$

Operator Ψ is *increasing* if and only if for all $A, B \subseteq X$,

$$A \subseteq B \Rightarrow \Psi(A) \subseteq \Psi(B)$$

Operator Ψ is *extensive* if and only $Y \subseteq \Psi(Y)$ for all $Y \subseteq X$, while Ψ is *antiextensive* if and only if $\Psi(Y) \subseteq Y$ for all $Y \subseteq X$.

Remark 8.19. Closure operators are increasing, idempotent, and extensive, while interior operators are increasing, idempotent, and antiextensive. Dilation is not idempotent: for $S \subseteq \mathbb{R}^d$ and $a, b \in \mathbb{R}^>$,

$$S \uparrow_a \uparrow_b = S \uparrow_{a+b}$$

In particular,

$$S \uparrow_a \uparrow_a = S \uparrow_{2a}$$

In contrast,

$$S \circ_a \circ_b = S \circ_{a \vee b}$$

Where $a \vee b$ is the maximum of a and b . When we perform a sequence of openings with varying radii, the result is the same as a single opening with the largest radius, so

$$S \circ_a \circ_a = S \circ_a$$

Further note that after a single application of an idempotent operator, further applications have no effect:

$$S \circ_a = S \circ_a \circ_a = S \circ_a \circ_a \circ_a = \dots$$

We identify sets unaffected by the application of an operator as elements of its invariance:

Definition 8.17. The *invariance* of an operator $\Psi : X \rightarrow X$, denoted $\mathcal{I}\Psi$, is defined as

$$\mathcal{I}\Psi := \{S \subseteq X \mid S = \Psi(S)\}$$

8.6.3. Self-duality and invariance with respect to set complement.

Remark 8.20. We now define self-duality and invariance with respect to set complement. Like duality, these properties describe the relationship between the way an operator transforms a set and the way it transforms the set's complement.

Definition 8.18. Operator Ψ is *self-dual* if and only if $\Psi(Y^c) = \Psi(Y)^c$.

Operator Ψ is *invariant with respect to set complement* if and only if $\Psi(Y) = \Psi(Y^c)$.

Remark 8.21. Examples of self-dual operators include Mason [67] and the morphological center [55]. Boundary operators are invariant with respect to set complement. The medial axis [15, 46, 70] is invariant with respect to set complement when it is defined as the singularities of the distance function to a set's boundary. However, that is not how we define the medial axis in Section 9.3.

8.7. Convex structures

Summary. We define convex structures, which capture the properties common to the Euclidean convex hull, relative convex hull, geodesic convex hull, and Euclidean

affine hull. We distinguish between convex geometries, whose hull operators behaves like the Euclidean convex hull, and matroids, whose hull operators behave like the Euclidean affine hull [63].

8.7.1. Convex structures.

Remark 8.22. Convex structures are differentiated from closure structures by the “stability under nested union” of the sets they contain. If a collection of convex sets is totally ordered by set inclusion, then the union of the sets in the collection is convex [63].

Definition 8.19. For a set X , suppose $\mathbb{C} \subseteq \wp(X)$. Then the pair (X, \mathbb{C}) is a *convex structure* if and only if for any $\mathbb{D} \subseteq \mathbb{C}$ such that the elements of \mathbb{D} are totally ordered by set inclusion, $\bigcup \mathbb{D} \in \mathbb{C}$.

If (X, \mathbb{C}) is a convex structure, we say the elements of \mathbb{C} are *convex*. The closure operator for a convex structure is a *convex hull operator*.

Remark 8.23. Euclidean convexity, relative convexity, geodesic convexity, and affinity are convex structures.

Not every closure structure is a convex structure. For example, the union of closed disks centered on the origin with radii in the interval $(0, 1)$ is an open disk. Because a union of nested closed disks is not necessarily closed, the closure structure defined by topologically closed subsets of \mathbb{R}^2 is not a convex structure.

In the limit as $r \rightarrow \infty$, morphological r -closure resembles the Euclidean convex hull. However, morphologically closed sets are always topologically closed, and as in our example involving unions of disks before, the union of topologically closed sets may be morphologically open. As a result, the analog of morphological closure using open halfspaces rather than open balls of finite radius does not define a convex structure.

8.7.2. Convex geometries and matroids.

Summary 6. We differentiate convex geometries and matroids from other convex structures based on whether their convex hull operators satisfy the exchange or antiexchange axiom. Convex geometries satisfy the antiexchange axiom and, like the Euclidean convex hull, add points that are “between” points currently in a set. Matroids satisfy the exchange axiom and, like the affine hull, add to a set all the points in the space it “spans.”

Definition 8.20. A convex structure (X, \mathbb{C}) with convex hull operator Ψ is a *convex geometry* if and only if for all $Y \subseteq X$ and all distinct $p, q \notin \Psi(Y)$,

$$q \in \Psi(Y \cup \{p\}) \Leftrightarrow p \notin \Psi(Y \cup \{q\}) \quad (1)$$

If instead

$$q \in \Psi(Y \cup \{p\}) \Leftrightarrow p \in \Psi(Y \cup \{q\}) \quad (2)$$

Then convex structure (X, \mathbb{C}) is a *matroid*. Property (1) is known as the *antiexchange axiom*, while (2) is the *exchange axiom* [63].

Remark 8.24. Euclidean convexity, relative convexity, and geodesic convexity are convex geometries. The Euclidean affinity is a matroid. The convex hull operator for a convex geometry formalizes “betweenness,” while the convex hull for a matroid formalizes “spanning.” For instance, if S is a connected subset of \mathbb{R}^2 and p is a points, the convex hull of $S \cup \{p\}$ adds to S every line segment connecting p to a point in S , which can be interpreted as the set of all points “between” S and p . The affine hull of $S \cup \{p\}$ has the same dimensionality as S if and only if p lies in the same affine subspace as S . Otherwise, $S \cup \{p\}$ “spans” a higher-dimensional space than S , and the dimensionality of the affine hull of $S \cup \{p\}$ is one greater than the dimensionality of the affine hull of S .

8.8. Regularity and symmetry

Summary. We provide definitions of open-regular sets and closed-regular sets, and we introduce new definitions of thin boundaries and irreducibility. We introduce topologically regular sets and then we redefine morphologically regular sets [1, 55] in terms of our new definition of regularity structures. We define symmetry with respect to set complement, and we characterize irreducible sets, regularity structures, and the medial cover as symmetric with respect to set complement.

8.8.1. Open-regular sets, closed-regular sets, and irreducibility.

Summary 7. Compared to arbitrary subsets of Euclidean space, open-regular and closed-regular sets have boundaries that are “thin” and “well-behaved.” An arbitrary set can be transformed into an open-regular or closed-regular set by applying a combination of interior and closure operations to it. However, further applications of interior and closure operators do not make an open-regular or closed-regular set’s boundary “thinner” or “better-behaved.” In particular, a composition of interior and closure operators cannot make a set’s boundary manifold.

Substituting morphological closure for topological closure, we obtain morphologically open-regular sets and morphologically closed-regular sets, whose behavior mirrors the behavior of topologically open-regular and closed-regular sets. Just as repeated applications of topological closure and interior are not guaranteed to produce a set with a manifold boundary, repeated applications of morphological opening or closing are not guaranteed to produce a set whose mortar has an empty interior.

We introduce a definition of irreducibility to formalize the concept of a set whose boundary as given by a particular closure structure’s boundary operator cannot be made smaller or thinner using the structure’s interior and closure operators. Open-regular and closed-regular sets are irreducible in the structure defined by topological closure, while morphologically open-regular and closed-regular sets are irreducible in the structure defined by morphological closure.

Definition 8.21. Let (X, A) be a closure structure with interior operator \circ_A , closure operator \bullet_A , and boundary operator ∂_A . A set $S \subseteq X$ has a *thin boundary* in (X, A) if and only if $(\partial_A S) \circ_A = \emptyset$.

Remark 8.25. Equivalently, a set S has a thin boundary in (X, A) if and only if $\partial_A S \in \mathcal{I}\partial_A$. Both open and closed sets have thin boundaries: $\partial(S\bullet_A) \subseteq \mathcal{I}\partial_A$ and $\partial(S\circ_A) \subseteq \mathcal{I}\partial_A$. There also exist closure structures containing sets that are neither open nor closed and yet have thin boundaries.

Definition 8.22. A set $S \subseteq \mathbb{R}^d$ is *open-regular* if and only if $S = S^{-\circ}$, and *closed-regular* if and only if $S = S^{\circ-}$.

Theorem 8.2. (Serra [55]) *Let (X, A) be a closure structure with closure operator \bullet_A and interior operator \circ_A . Then the compositions $\circ_A\bullet_A$ and $\bullet_A\circ_A$ are idempotent.*

Remark 8.26. Because either composition of any closure structure's closure and interior operator are idempotent, the compositions \circ^- and $^- \circ$ are idempotent in particular. Writing $S^{\circ-} = S^{\circ-\circ-}$, we see $S^{\circ-}$ is closed-regular for all $S \subseteq \mathbb{R}^d$. Similarly, $S^{-\circ} = S^{-\circ-\circ}$, so $S^{-\circ}$ is open-regular for all $S \subseteq \mathbb{R}^d$.

Because \circ , $-$, \circ^- , and $^- \circ$ are idempotent, the only distinct compositions of \circ and $-$ are \circ , $-$, \circ^- , $^- \circ$, $\circ-\circ$, and $^- \circ^-$. For all $S \subseteq \mathbb{R}^d$, the sets $S^{\circ-}$ and $S^{-\circ-}$ are closed-regular, while $S^{-\circ}$ and $S^{\circ-\circ}$ are open-regular.

Example 8.2. The boundaries of open-regular and closed-regular sets have properties that make them suitable for many solid modeling tasks. However, their boundaries are not necessarily manifold, which is a condition required by many geometric algorithms. For example, no composition of closure and interior operators makes the boundary of the set shown in Figure 8.1 manifold.

Definition 8.23. A set $S \subseteq \mathbb{R}^d$ is *morphologically open r -regular* if and only if $S = S \bullet_r \circ_r$, and *morphologically closed r -regular* if and only if $S = S \circ_r \bullet_r$.

FIGURE 8.1. On the left, we show a set S in black consisting of two squares rotated by $\frac{\pi}{2}$ and positioned so their corners touch. On the right we draw the boundary of S in red. The closure of S consists of a single connected component, while the interior of S consists of two. No matter which subset of its boundary S includes, the boundary of S remains nonmanifold.

Remark 8.27. Because $\circ_r \bullet_r$ and $\bullet_r \circ_r$ are idempotent, $S \circ_r \bullet_r = S \circ_r \bullet_r \circ_r \bullet_r$, so $S \circ_r \bullet_r$ is closed-regular for all $S \subseteq \mathbb{R}^d$. Similarly, $S \bullet_r \circ_r = S \bullet_r \circ_r \bullet_r \circ_r$, so $S \bullet_r \circ_r$ is open-regular for all $S \subseteq \mathbb{R}^d$.

Because \circ_r , \bullet_r , $\circ_r \bullet_r$ and $\bullet_r \circ_r$ are idempotent, the only distinct compositions of \circ_r and \bullet_r are \circ_r , \bullet_r , $\circ_r \bullet_r$, $\bullet_r \circ_r$, $\circ_r \bullet_r \circ_r$, and $\bullet_r \circ_r \bullet_r$. For all $S \subseteq \mathbb{R}^d$, the sets $S \circ_r \bullet_r$ and $S \bullet_r \circ_r \bullet_r$ are closed- r -regular, while $S \bullet_r \circ_r$ and $S \circ_r \bullet_r \circ_r$ are open- r -regular.

Definition 8.24. Let (X, A) be a closure structure with interior operator \circ_A , closure operator \bullet_A , and boundary operator ∂_A . A set $S \subseteq X$ is *irreducible* in (X, A) if and only if $\partial_A(S \circ_A) = \partial_A S = \partial_A(S \bullet_A)$.

Remark 8.28. Open-regular and closed-regular sets are irreducible in $(\mathbb{R}^d, \mathcal{I}^-)$. Morphologically open r -regular and closed r -regular sets are irreducible in $(\mathbb{R}^d, \mathcal{I}_{\bullet_r})$. Irreducible sets have thin boundaries, but their boundaries are not necessarily manifold.

The concepts of irreducibility and thin boundary both concern the “niceness” or “simplicity” of a set’s boundary in a given closure structure. Thin boundaries have empty interiors, but there are thin boundaries that we can make smaller by applying closure and interior operators. If S is irreducible in (X, A) , by contrast, its boundary is invariant under its closure structures’ interior, closure, and boundary operators: $\partial_A S \in \mathcal{I} \circ_A$, $\partial_A S \in \mathcal{I} \bullet_A$, and $\partial_A S \in \mathcal{I} \partial_A$. Note, however, that a set that is irreducible in one closure is not necessarily irreducible in another. An irreducible set S in $(\mathbb{R}^d, \mathcal{I}_{\bullet_r})$, for instance, may have a mortar with a nonempty topological interior: although $(\mathcal{M}_r S) \circ_r = \emptyset$, we have $(\mathcal{M}_r S)^\circ \neq \emptyset$.

8.8.2. Topologically and morphologically regular sets.

Summary 8. We introduce a new definition of topologically regular sets, which have manifold boundaries separating their interiors and exteriors. While the definition of topological regularity does not specify which portion of its boundary a topologically regular set contains, the closure of a topologically regular set is closed-regular, while the interior of a topologically regular set is open-regular. A closed-regular or open-regular set, by contrast, may not have a manifold boundary or a boundary that can be made manifold by the application of closure and interior operators.

We also introduce a definition of the regularity structure generated by an interior operator, which consists of topologically regular sets whose topological interior and exterior consist of connected components invariant under the given interior operator. Specifically, we use regularity structures to redefine morphologically r -regular sets [1, 55], whose interior and exterior consists of connected components invariant under morphological r -opening. We outline geometric properties of morphologically regular sets, such as that the sectional curvatures of a morphologically r -regular set lie in $[-\frac{1}{r}, \frac{1}{r}]$, making them qualitatively smooth and blobby.

Definition 8.25. A set $S \subseteq \mathbb{R}^d$ is *topologically regular* if and only for every $p \in \partial S$ there exists an $r \in \mathbb{R}^>$ such that for all $\rho \in (0, r)$, both $B_\rho(p) \cap S^\circ$ and $B_\rho(p) \cap S^{c\circ}$ are homeomorphic to an open ball.

Remark 8.29. If $S \subseteq \mathbb{R}^d$ is topologically regular, ∂S is a $(d - 1)$ -manifold separating S from S^c . Topologically regular sets are irreducible in $(\mathbb{R}^d, \mathcal{I}^-)$.

Definition 8.26. Let (\mathbb{R}^d, A) be a closure structure with interior operator Ψ . Then $S \subseteq \mathbb{R}^d$ is an element of the *regularity structure* generated by Ψ if and only if

- (1) S is topologically regular.
- (2) For every $K \in \kappa(\mathbb{R}^d \setminus \partial S)$, we have $K \in \mathcal{I}\Psi$.

Definition 8.27. A set $S \subseteq \mathbb{R}^d$ is *morphologically r -regular* if and only if it is an element of the regularity structure generated by \circ_r . If S is not morphologically r -regular, we de r -irregular.

Remark 8.30. Membership in a regularity structure is a stronger condition than irreducibility: if (\mathbb{R}^d, A) is a closure structure with interior operator Ψ , then an element of the regularity structure generated by Ψ is irreducible in (\mathbb{R}^d, A) , but an irreducible set in (\mathbb{R}^d, A) may not be part of the regularity structure generated by Ψ . For instance, if $S \subseteq \mathbb{R}^d$ is morphologically r -regular, then it is irreducible in $(\mathbb{R}^d, \mathcal{I}_{\bullet, r})$. However, there exists a set $S \subseteq \mathbb{R}^d$ irreducible in $(\mathbb{R}^d, \mathcal{I}_{\bullet, r})$ such that $(\mathcal{M}_r S)^\circ \neq \emptyset$, implying it is not topologically regular.

Remark 8.31. If S is morphologically regular, for every ball $b \in \mathbb{B}_{2r}(\partial S)$ of radius $2r$ whose center lies on the boundary of S , both $b \cap S^\circ$ and $b \cap S^{c\circ}$ are homeomorphic to an open ball. Moreover, every point $p \in \partial S$ lies on both the boundary of an open ball of radius r contained in S and the boundary of an open ball of radius r contained in S^c . Under Definition 10.20, all sectional curvatures of S at p lie in $[-\frac{1}{r}, \frac{1}{r}]$.

8.8.3. Symmetry with respect to set complement.

Summary 9. We define a collection of sets as symmetric with respect to set complement if and only if whenever the collection contains a set S , it also contains S^c . We then define an operator that maps a pair of sets to a third set as symmetric with respect to set complement if swapping the order of the sets in the pair exchanges the interior and exterior of the operator's output but preserves its boundary.

Definition 8.28. For a set X , let $\mathbb{C} \subseteq \wp(X)$ be a collection of sets. Then \mathbb{C} is *symmetric with respect to set complement* if and only if for all $Y \subseteq X$,

$$Y \in \mathbb{C} \Leftrightarrow Y^c \in \mathbb{C}$$

Remark 8.32. A regularity structure is symmetric with respect to set complement, as is the collection of sets irreducible in a given closure structure. Closed sets generally are not symmetric with respect to set complement. For instance, neither topologically closed nor morphologically r -closed sets are symmetric with respect to set complement.

Instead, the complement of a closed set in a closure structure is open, and the closure operator in a closure structure pairs with a dual interior operator. Collections of sets that behave symmetrically with respect to set complement lead us to define operators that, unlike dual operators, behave symmetrically with respect to set complement.

Definition 8.29. For a set X and $A, B \subseteq X$, operator $\Psi : \wp(X) \times \wp(X) \rightarrow \wp(X)$ is *symmetric with respect to set complement* if and only if

$$\Psi(A, B)^{\circ} = \Psi(B, A)^{c\circ} \wedge \Psi(A, B)^{-} = \Psi(B, A)^{c-}$$

Remark 8.33. The medial cover of R relative to G defined in Chapter 15 is symmetric with respect to set complement under weak assumptions on R and G . The tight hull is not symmetric with respect to set complement under this definition because the tight hull of R relative to G is not necessarily unique.

CHAPTER 9

THE MEDIAL AXIS

9.1. Overview

We begin by defining the distance transform, which maps every point in space to its distance from a given set. We then introduce the regularity transform, which maps every point in space to the radius of the largest open ball containing the point that lies within a given set. We define maximal balls, then define the medial axis in terms of maximal balls (Section 9.3.) We express interior, opening, and dilation operators in terms of maximal balls. We introduce a bloom operator that maps a point on the medial axis to the maximal ball centered at that point. We define the contact set of a maximal ball b in set $S \subseteq \mathbb{R}^d$ as the intersection of ∂b with ∂S (Definition 9.11.) We then introduce a comb operator that maps a point of the medial axis to the convex hull of its contact set. We use the comb operator to reformulate the definition of the α -shape [19, 20], and we define β -shape as the dual of the α -shape. We conclude with a brief explanation of the piecewise real analytic boundary assumption used by Choi, Choi, and Moon [15]. Any set satisfying that assumption is homeomorphic to its medial axis. If the set is two-dimensional, its medial axis has a graph structure consisting of a finite number of disjoint curved edges meeting at a finite number of vertices.

9.2. Distance and regularity transforms

Summary 10. We define the distance transform, then introduce a novel definition of the regularity transform. We show that we can obtain dilation from the distance transform and morphological opening from the regularity transform. We observe that

the maximal balls that determine the regularity transform of a set are centered on singular points of the distance transform of the set's complement.

9.2.1. Distance transform.

Definition 9.1. For $S \subseteq \mathbb{R}^d$, the *distance transform* $\mathcal{D}^S : \mathbb{R}^d \rightarrow \mathbb{R}^{\geq}$ is defined as

$$\mathcal{D}^S(p) := \bigwedge \{r \in \mathbb{R}^{\geq} \mid p \in S \uparrow_r^-\}$$

Remark 9.1. This defines the distance of a point $p \in \mathbb{R}^d$ to a set $S \subseteq \mathbb{R}^d$ as the least radius value $r \in \mathbb{R}^{\geq}$ such that p lies in the closure of the dilation of S by r . The distance transform specializes the Euclidean distance function:

$$\mathcal{D}^S(p) = d(p, S) = d(S, p)$$

Remark 9.2. We can obtain dilation by thresholding the distance transform:

$$S \uparrow_r = \{p \in \mathbb{R}^d \mid \mathcal{D}^S(p) < r\}$$

This follows directly from the definition of the distance transform. Paraphrasing, “The dilation of S by r is equal the set of all points contained in the interior of the closure of a dilation of S by r or less.”

9.2.2. Regularity transform.

Summary 11. We introduce the regularity transform using the same mechanisms that define the distance transform. The relationship between the distance transform and the regularity transform is the relationship between distance and size, or the relationship between a line segment of a given length and a disk of a given radius. Informally, the regularity of a point in a set is the radius of the largest disk in the set containing the point. It measures the “thickness” of the set around the point, giving it close relationships to existing concepts like least and local feature size. As indicated

in previous work [18, 55], measures of thickness can be used to determine the rate at which a solid can be sampled so that its reconstruction from its samples has an error bounded by the intersample spacing.

Definition 9.2. For $S \subseteq \mathbb{R}^d$, the *regularity transform* $\mathcal{R}^S : \mathbb{R}^d \rightarrow \mathbb{R}^{\geq}$ is defined as

$$\mathcal{R}^S(p) := \bigvee \{ \{r \in \mathbb{R}^> \mid p \in S \circ_r^-\} \cup 0 \}$$

We refer to $\mathcal{R}^S(p)$ as the *regularity* of p with respect to S .

The *regularity* of S is defined as $\bigvee_{p \in S^\circ} \mathcal{R}^S(p)$.

Claim 9.1. The regularity of point $p \in \mathbb{R}^d$ with respect to $S \subseteq \mathbb{R}^d$ is the greatest radius r such that p lies in the closure of the opening of S by r . Equivalently, the regularity of p with respect to S is the radius of the largest open ball in S whose closure contains p :

$$\mathcal{R}^S(p) = \bigvee \{ \{r \in \mathbb{R}^> \mid \exists b \in \mathbb{B}_r^S : p \in b^-\} \cup 0 \}$$

Remark 9.3. We can obtain opening by thresholding the regularity transform:

$$S \circ_r = \{p \in \mathbb{R}^d \mid \mathcal{R}^S(p) \geq r\}^\circ$$

As with the distance transform, this follows from the definition of the regularity transform. Paraphrasing, “The opening of S by r is the interior the set of all points contained in the closure of the opening of S by r or less.”

Remark 9.4. For $S \subseteq \mathbb{R}^d$ and $p \in S^\circ$, suppose $b \in \mathbb{B}^S$ is a ball of maximal radius containing p . Then the center of b lies on a singular point of the distance transform of S^c , where its gradient is undefined. Using the terminology of Section 9.3, b is a maximal ball centered on the medial axis of S . Using the terminology of Section 9.4, we obtain the regularity transform with a variant of the bloom of the medial axis.

9.2.3. Transforms from increasing and decreasing operators.

Summary 12. The distance and regularity transforms suggest a method for generating transforms from operators parameterized by scale that either progressively increase or decrease the size of a set (as determined by the ordering of sets under set inclusion) when the scale parameter varies monotonically. We define monotonically increasing operators, which increase the size of a set as a scale parameter increases, and monotonically decreasing operators, which decrease set size as scale increases. We then introduce a novel definition of expansion transforms in terms of monotonically increasing operators, characterizing the distance transform as an expansion transform. We also introduce contraction transforms, which we define in terms of monotonically decreasing operators. We characterize the regularity transform as a contraction transform. We conclude by defining total dual transforms, which combine an expansion transform generated by a given operator with a contraction transform generated by the operator's dual.

Definition 9.3. Let $\Psi : \mathbb{R}^> \times \wp(\mathbb{R}^d) \rightarrow \wp(\mathbb{R}^d)$ be an operator parameterized by a positive real value $r \in \mathbb{R}^>$ that maps one subset of Euclidean space to another. Then Ψ is a *monotonically increasing operator* if and only if for all $S \subseteq \mathbb{R}^d$ and $a, b \in \mathbb{R}^>$,

$$a \leq b \Leftrightarrow \Psi_a(S) \subseteq \Psi_b(S)$$

Example 9.1. Dilation and closing are monotonically increasing operators.

Definition 9.4. Let $\Psi : \mathbb{R}^> \times \wp(\mathbb{R}^d) \rightarrow \wp(\mathbb{R}^d)$ be an operator parameterized by a positive real value $r \in \mathbb{R}^>$ that maps one subset of Euclidean space to another. Then Ψ is a *monotonically decreasing operator* if and only if for all $S \subseteq \mathbb{R}^d$ and $a, b \in \mathbb{R}^>$,

$$a \leq b \Leftrightarrow \Psi_a(S) \supseteq \Psi_b(S)$$

Example 9.2. Erosion and opening are monotonically decreasing operators.

Definition 9.5. For $S \subseteq \mathbb{R}^d$ and monotonically increasing operator Ψ , assume $\lim_{r \rightarrow 0} \Psi_r(S) = S^-$ and $\lim_{r \rightarrow \infty} \Psi_r(S) = \mathbb{R}^d$. Then the *expansion transform* $\mathcal{E}_\Psi^S :$

$\mathbb{R}^d \rightarrow \mathbb{R}^{\geq}$ of S under Ψ is defined as

$$\mathcal{E}_{\Psi}^S(p) := \bigwedge \{r \in \mathbb{R}^> \mid p \in \Psi_r(S)^-\}$$

Example 9.3. The expansion transform of $S \subseteq \mathbb{R}^d$ under dilation is the distance transform. The expansion transform of S under morphological closing is the regularity transform of S^c .

Definition 9.6. For $S \subseteq \mathbb{R}^d$ and monotonically decreasing operator Ψ , assume $\lim_{r \rightarrow 0} \Psi_r(S) = S^\circ$ and $\lim_{r \rightarrow \infty} \Psi_r(S) = \emptyset$. The *contraction transform* $\mathcal{C}_{\Psi}^S : \mathbb{R}^d \rightarrow \mathbb{R}^{\geq}$ of S under Ψ is defined as

$$\mathcal{C}_{\Psi}^S(p) := \bigvee \{r \in \mathbb{R}^> \mid p \in \Psi_r(S)^-\}$$

Example 9.4. The contraction transform of $S \subseteq \mathbb{R}^d$ under morphological opening is the regularity transform. The contraction transform of S under erosion is the distance transform of S^c .

Definition 9.7. Let Ψ be a monotonically increasing operator, and let Υ be the dual of Ψ . The total transform of $S \subseteq \mathbb{R}^d$ under Ψ at $p \in \mathbb{R}^d$, denoted $\mathcal{T}_{\Psi}^S(p)$, is defined as

$$\mathcal{T}_{\Psi}^S(p) := \begin{cases} \mathcal{E}_{\Psi}^S(p) & p \in S^{\circ\circ} \\ \mathcal{C}_{\Upsilon}^S(p) & p \in S^\circ \\ \mathcal{E}_{\Psi}^S(p) \wedge \mathcal{C}_{\Upsilon}^S(p) & p \in \partial\partial S \\ 0 & (\partial S)^\circ \end{cases}$$

Example 9.5. For $S \subseteq \mathbb{R}^d$, the total transform \mathcal{T}_{\uparrow}^S returns the distance of a point $p \in \mathbb{R}^d$ to ∂S :

$$\mathcal{T}_{\uparrow}^S(p) = d(p, \partial S)$$

For $p \in S^\circ$, the total transform $\mathcal{T}_\bullet^S(p)$ is equal to the value of the regularity transform $\mathcal{R}^S(p)$, while for $p \in S^{c^\circ}$ we have $\mathcal{T}_\bullet^S(p) = \mathcal{R}^{S^c}(p)$. For $p \in (\partial S)^\circ$, we have $\mathcal{T}_\bullet^S(p) = 0$. However, for some $S \subseteq \mathbb{R}^d$ there may exist a point $p \in \partial\partial S$ such that the regularity of p with respect to S and the regularity of p with respect to S^c are nonzero and nonequal. In this case, the value of $\mathcal{T}_\bullet^S(p)$ is the minimum of the two regularity values. For instance, if S is an open disk of radius r and $p \in \partial S$, we have $\mathcal{T}_\bullet^S(p) = r$.

9.3. Maximal balls and the medial axis

Summary 13. Given a subset S of Euclidean space, we define the set of maximal balls in S . We then define the medial axis of S as the set of the centers of the maximal balls in S . We show that topological opening, morphological opening, and dilation can be expressed in terms of maximal balls.

Definition 9.8. A ball $b \in \mathbb{B}^S$ is *maximal* in $S \subseteq \mathbb{R}^d$ if and only if no ball $a \in \mathbb{B}^S$ exists such that $b \subset a$.

Notation 4. Let \mathbb{M}^S denote the set of all balls maximal in $S \subseteq \mathbb{R}^d$:

$$\mathbb{M}^S = \{b \in \mathbb{B}^S \mid \sim \exists a \in \mathbb{B}^S : b \subset a\}$$

Definition 9.9. The *medial axis* of $S \subseteq \mathbb{R}^d$, denoted MA^S , is defined as

$$\text{MA}^S := \{\pi(b) \mid b \in \mathbb{M}^S\}$$

Remark 9.5. We can express interior, opening, and erosion using only the maximal balls in a set. We contrast this to Definitions 8.2, 8.3, and 8.5, which use all the open balls in a set.

$$S^\circ = \bigcup \mathbb{M}^S$$

$$S_{\circ_r} = \bigcup \mathbb{M}_{\geq r}^S$$

$$S \uparrow_r = \bigcup_{b \in \mathbb{M}^S} b \uparrow_r$$

Topological closure, morphological closure, and dilation can then be expressed as dual operators: $- = {}^{c \circ c}$, $\bullet_r = {}^c \circ_r {}^c$, and $\uparrow_r = {}^c \downarrow_r {}^c$.

9.4. Bloom, contact set, and radius function

Summary 14. We introduce a definition of the bloom of a subset of the medial axis as the union of the maximal balls centered on that subset. We then define the contact set of a maximal ball in a set $S \subseteq \mathbb{R}^d$ as the intersection of its boundary with ∂S . We define the radius function for the medial axis, and we illustrate filtering the medial axis based on its radius function.

Definition 9.10. For $S \subseteq \mathbb{R}^d$, the *bloom* $\text{BL}^S : \wp(\text{MA}^S) \rightarrow \wp(\mathbb{R}^d)$ is defined as

$$\text{BL}^S(X) := \bigcup \mathbb{M}^S(X)$$

For $p \in \mathbb{R}^d$, let $\text{BL}^S(p) := \text{BL}^S(\{p\})$.

Remark 9.6. The bloom of the medial axis of $S \subseteq \mathbb{R}^d$ is equal to the interior of S : $\text{BL}^S(\text{MA}^S) = S^\circ$. The bloom operator is an “inverse medial axis transform:” it maps a point on the medial axis to the maximal disk centered on it.

Definition 9.11. For $S \subseteq \mathbb{R}^d$ with $b \in \mathbb{M}^S$, the *contact set* $\text{CS}^S(b)$ is defined as $\text{CS}^S(b) := \partial b \cap \partial S$.

The contact set of $X \subseteq \text{MA}^S$ is defined as $\text{CS}^S(X) := \text{BL}^S(X) \cap \partial S$.

The contact set of $p \in \text{MA}^S$ is defined as $\text{CS}^S(p) := \text{CS}^S(\{p\})$.

The set of *contact components* of a contact set C is defined as $\kappa(C)$.

Remark 9.7. A maximal ball has at least one contact component. Under modest assumptions, almost all maximal balls have two contact components, and no maximal

ball has more than $d + 1$ contact components in \mathbb{R}^d . The number of contact components of a maximal ball relates to the graph structure of the medial axis, as discussed in Section 9.6.

Definition 9.12. For $S \subseteq \mathbb{R}^d$, the *radius function* $\rho^S : \text{MA}^S \rightarrow \mathbb{R}$ maps points of the medial axis of S to the radii of the maximal balls centered on them:

$$\rho^S(p) := \rho(\text{BL}^S(p))$$

Notation 5. For $S \subseteq \mathbb{R}^d$ and $r \in \mathbb{R}^>$ with binary relation \sim on \mathbb{R} , let $\text{MA}_{\sim r}^S$ denote the set of points $p \in \text{MA}^S$ such that $\rho^S(p) \sim r$:

$$\text{MA}_{\sim r}^S := \{p \in \text{MA}^S \mid \rho^S(p) \sim r\}$$

Remark 9.8. We can express opening in terms of the bloom of a trimmed medial axis:

$$S \circ_r = \text{BL}^S(\text{MA}_{\geq r}^S)$$

Similarly, erosion trims the medial axis. If S is open,

$$\text{MA}^{S \downarrow r} = \text{MA}_{> r}^S$$

While if S is closed,

$$\text{MA}^{S \downarrow r} = \text{MA}_{\geq r}^S$$

Dilation and morphological closure may both add and remove points from a set's medial axis. By duality, erosion and morphological opening may both add and remove to the medial axis of a set's complement.

9.5. Comb and alpha shape

Summary 15. We introduce the comb operator, which maps a maximal disk to the convex hull of its contact set. Applied to a set $S \subseteq \mathbb{R}^2$ where no maximal disk has more than three contact points, the comb of the medial axis of S decomposes it into simplices.

Definition 9.13. For $S \subseteq \mathbb{R}^d$, the *comb* $\text{CM}^S : \wp(\text{MA}^S) \rightarrow \wp(\mathbb{R}^d)$ is defined as the union over $p \in X$ of the convex hull of the contact set of p :

$$\text{CM}^S(X) := \bigcup_{p \in X} \text{CH}(\text{CS}^S(p))$$

Remark 9.9. The comb of the medial axis of $S \subseteq \mathbb{R}^d$ is equal to the closure of its interior: $\text{CM}_S(\text{MA}_S) = S^{\circ-}$.

Definition 9.14. The β -*shape* of $S \subseteq \mathbb{R}^d$, denoted \mathcal{B}_β^S , is defined as

$$\mathcal{B}_\beta^S := \text{CM}^S(\text{MA}_{\leq \beta}^S)$$

The α -*shape* of S , denoted \mathcal{A}_α^S , is defined as $\mathcal{A}_\alpha^S := \mathcal{B}_\alpha^{S^c}$.

Remark 9.10. The β -shape is a polygonal analog of opening, where the disks in opening are replaced by convex hulls of contact points of disks with radii greater than or equal to β . Our definition of the α -shape of $S \subseteq \mathbb{R}^d$ is consistent with Edelsbrunner's [19, 20] when S consists of a finite number of points. It may contain simplices ranging in dimension from 0 to d . Applied to other kinds of sets, our α -shape performs a polygonal analog of closing. The contraction transform under \mathcal{B} is analogous to the regularity transform, and the total transform $\mathcal{T}_\mathcal{A}^S$ is analogous to \mathcal{T}_\bullet^S .

9.6. Piecewise real analytic boundary assumption

Summary 16. The structure of a set's medial axis can be complex if the set's boundary is complex. For instance, if a set is bounded by a fractal or has oscillatory behavior like that exhibited by $\sin \frac{1}{x}$, the boundary of the set's medial axis may have a nonempty interior. Choi, Choi, and Moon [15] demonstrate that if the boundary of a two-dimensional set is a manifold equal to the union of a finite number of real analytic pieces, its medial axis has a finite graph structure. The assumption is sufficiently nonrestrictive to be applicable to a useful class of sets encountered in solid modeling.

Below, we simplify and adapt terminology and theorems from Choi, Choi, and Moon. Their paper and related works by Wolter [70, 71] and Siddiqi and Pizer [46] provide greater detail.

Remark 9.11. Informally, a function $f : \mathbb{R} \rightarrow \mathbb{R}$ of one variable is real analytic if and only if for every $x \in \mathbb{R}$, the Taylor series of f about x converges to f in the neighborhood of x . A nonzero real analytic function can only have a finite number of zeros on a bounded domain. Because the derivatives of a real analytic function are also real analytic functions, a real analytic function has a finite number of maxima, minima, and inflection points.

Suppose a segment $X \subseteq \mathbb{R}^2$ of a set's boundary is a finite-length curve homeomorphic to a line segment. When we describe X as a real analytic curve, we assert that its coordinates are real analytic functions of its arclength. Its curvature is then a real analytic function and so has a finite number of critical points. The distance function $g^X : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined by $g^X(p) := d(p, X)$ is also real analytic, and if we construct distance functions for a finite number of boundary segments, the sums and differences of the distance functions are real analytic, with real analytic zero level sets. A medial axis point is generally equidistant from either two points on the same boundary segment or one point on each of two different segments. Medial axis segments then lie on a segment from the zero level set of one boundary segment's distance function

or on a segment from the zero level set of the difference of two distance functions. In either case, a medial axis segment is a real analytic curve. Moreover, the segment's radius function (Definition 9.12) is real analytic.

Broadly, the assumption that the boundary of a set consists of a finite number of real analytic pieces ensures that the “complexity” of its medial axis is finite, because the functions that characterize its behavior have a finite number of critical points.

Assumption 1. *For a bounded, topologically regular set $S \subseteq \mathbb{R}^2$, assume ∂S is equal to the union of a finite number of real analytic segments.*

Remark 9.12. If a set satisfies Assumption 1, its medial axis consists of a finite number of edges, which may be curved, joining at a finite number of vertices. If a set does not satisfy Assumption 1, it may have an infinite number of edges, and an infinite number of edges may meet at a vertex.

Whether a medial axis point is part of a vertex or lies on an edge depends on its number of contact components, which we define as its degree:

Definition 9.15. The *degree* of $p \in \text{MA}^S$, denoted $\text{deg}^S(p)$, is the cardinality of its set of contact components:

$$\text{deg}^S(p) := \#\kappa(\text{CS}^S(p))$$

Remark 9.13. At a point $p \in \text{MA}^S$,

$$\text{deg}^S(p) = \lim_{r \rightarrow 0} \#\kappa((B_r(p) \cap \text{MA}^S) \setminus p)$$

The degree of a medial axis point is the number of connected components that remain in its neighborhood when the point is removed.

Definition 9.16. The set of *edges* of the medial axis of $S \subseteq \mathbb{R}^d$, denoted E^S , is defined as

$$E^S := \kappa(\{p \in MA^S \mid \deg^S(p) = 2\})$$

Remark 9.14. Although under special circumstances an edge can be a simple closed curve, in most cases of interest to us an edge is homeomorphic to an open line segment. Assumption 1 guarantees that edges are real analytic curves and not fractal or otherwise pathological curves from the perspective of solid modeling. By definition, edges do not intersect. They meet only at vertices located at their endpoints.

Definition 9.17. The set of *vertices* of the medial axis of S , denoted V^S , is defined as

$$V^S := \bigcup_{e \in E^S} e^- \setminus e$$

Remark 9.15. The closure of an edge includes the points at the edge’s endpoints. Because the edges themselves do not include their endpoints, the vertices of the medial axis are precisely the edge endpoints.

Remark 9.16. A generic medial axis point has two contact components and lies on an edge. In nonsingular situations in two dimensions, a medial axis point with more than two contact components is a vertex with exactly three contact components. While a medial axis point with two contact components lies on a singular ridge of the distance function to the boundary of a set, a medial axis point with three contact components lies at the intersection of two singular ridges. A medial axis point may also have a single contact component, in which case it is the tip of an edge, and the tip’s contact component is a circular arc or the tip is a center of curvature. While a medial axis point may have more than three contact components, this corresponds to an input set that is not in “general position.” By slightly perturbing the set, we can reduce the maximum number of medial axis contact components to three.

Theorem 9.1. (*Choi, Choi, and Moon [15]*) *Under Assumption 1, the medial axis consists of a finite number of edges and vertices.*

Remark 9.17. Informally, an endpoint of the medial axis of a set $S \subseteq \mathbb{R}^d$ arises from a local curvature maximum on ∂S . Curvature maxima are a kind of critical point, and under Assumption 1, there are a finite number of such points along ∂S . Because each edge has at most two endpoints, the number of medial axis vertices is also finite.

###

Theorem 9.2. (*Choi, Choi, and Moon*) Under Assumption 1, the edges and vertices of the medial axis define a finite graph structure.

Remark 9.18. The essence of the proof is that edges are disjoint and each edge connects at most two vertices.

Theorem 9.3. (*Choi, Choi, and Moon; Lieutier [41]*) Under Assumption 1, set S is homotopic to its medial axis.

Remark 9.19. In two dimensions, we can retract a set $S \subseteq \mathbb{R}^d$ onto its medial axis by moving each point $p \in \partial S$ along the line segment connecting p to the center of the maximal disk touching it at a rate proportional to the segment's length. With some refinement, the result holds when more than one maximal disk touches p . Lieutier [41] extends the result to higher dimensions.

If $S \subseteq \mathbb{R}^2$ is homotopic to its medial axis, then MA^S is connected if and only if S is connected, and the number of cycles in MA^S is equal to the number of holes in S . Retraction-based approaches to establishing the homotopy implicitly parameterize the interior of S , and such a parameterization may have uses beyond the proof of homotopic equivalence.

CHAPTER 10

NORMALS AND CURVATURES

10.1. Overview

This chapter defines slack, which plays a central role in our definition of tight hulls. To define slack, we introduce a novel definition of a normal to a subset of Euclidean space. We also define several means of describing normal variation for sets with nondifferentiable normal fields, including generalized forms of Gaussian and mean curvature, the Lipschitz continuity of the normals to a set, and the smooth offset interval. These measures prove important in our later discussion of the geometry of tightenings and tight hulls, whose boundaries generally do not have differentiable normal fields.

We begin by introducing material necessary for our definitions of normals and normal behavior. We informally describe the Hausdorff measure, which unifies measures of set size such as length, area, and volume. We formally define the geodesic distance between two points in a set, which is the length of the shortest path connecting them, and we use geodesic distance to define geodesic balls and geodesic neighborhoods of sets. We then define the dimensionality of a set, which indicates whether it consists of points, curves, solids, or higher-dimensional components.

We introduce a new approach to the definition of normals to arbitrary subsets of Euclidean space. Our definition proceeds in three steps. First, we define normals to morphologically r -closed sets. We then incorporate a limiting process into this definition to obtain normals to topologically closed sets. Finally, we define normals to arbitrary sets by considering the connected components of the set in the neighborhood of a point. This allows us to assign normal fields to sets that include fractals, point clouds, and other “complex” or “pathological” subsets. We define smooth and singular

boundary points, and we introduce a definition of normal count, which captures the idea of a normal being “used” more than once in certain singular situations.

We define measures of normal variation, beginning with slack. Slack is a measure of total normal variation closely related to the integral of the absolute value of Gaussian curvature, but its formulation in terms of our definitions of normals and normal behavior makes it behave differently from absolute Gaussian curvature in some cases. We introduce a sign function that reflects whether a point’s neighborhood eventually reverses its orientation when it is offset along its normals. We then modify our definition of slack using the sign function to obtain a new formulation of the definition of Gaussian curvature. We discuss Gaussian curvature on triangle meshes, and we describe how it concentrates at mesh vertices. We then introduce a form of normal offsetting different from dilation in order to define a new formulation of mean curvature, and we describe how it concentrates at mesh edges. We use our mean curvature definition to define the sectional curvatures through a normal at a point. From the interval of sectional curvatures we extract the maximum and minimum sectional curvature, which on a smooth surface correspond to the principal curvatures.

We apply the definition of Lipschitz continuity to a set’s normals, which lets us express a bound on the rate of change in normal direction over the set’s boundary. Finally, we provide a new definition of the smooth offset interval at a point, which characterizes how far we can offset the neighborhood of a point before it develops a singularity. We relate the smooth offset interval to morphological regularity and the sectional curvature interval.

10.2. Hausdorff measure, geodesic balls, and dimensionality

10.2.1. Hausdorff measure.

Remark 10.1. The d -dimensional Hausdorff measure of a set like a ball or cube in \mathbb{R}^d is its area, volume, etc., depending on the value of d . Hausdorff defined the measure 1918 [27], and Federer [23] offers an English-language presentation of it.

Unlike similar measures, such as the Lebesgue measure, the Hausdorff measure can measure lower-dimensional subsets of Euclidean space. The Lebesgue measure and the d -dimensional Hausdorff measure of a unit ball's boundary is zero, but the $(d - 1)$ -dimensional Hausdorff measure of the ball's boundary is the measure of the unit $(d - 1)$ -dimensional sphere, assuming the Hausdorff measure is scaled to agree with the Lebesgue measure. In general, the dimension d in a d -dimensional Hausdorff measure need not be an integer. We take d to be a nonnegative real.

Notation 6. Let H^d denote the d -dimensional Hausdorff measure scaled so that if X is a d -dimensional cube with unit edge length, $H^d(X) = 1$.

10.2.2. Geodesic distances, balls, and neighborhoods.

Remark 10.2. The geodesic distance between two points in a path-connected set $S \subseteq \mathbb{R}^d$ is the length (one-dimensional Hausdorff measure) of the shortest path connecting them that lies within S . The geodesic distance defines geodesic balls and spheres analogous the balls and spheres defined by the Euclidean metric. The geodesic neighborhood of a point $p \in \partial S$, which is a small geodesic ball in ∂S centered on p , proves important in many of our local measures of normal behavior.

Definition 10.1. Let $S \subseteq \mathbb{R}^d$ be path-connected, and for $p, q \in S$ let Π denote the set of all path-connected subsets of S containing both p and q . Then the *geodesic distance* between p and q in S , denoted $g^S(p, q)$, is defined as

$$g^S(p, q) := \bigwedge_{\pi \in \Pi} H^1(\pi)$$

For $p \in S$ and $X \subseteq S$, the geodesic distance between p and X in S , denoted $g^S(p, X)$, is defined as

$$g^S(p, X) := \bigwedge_{q \in X} g^S(p, q)$$

Let $g^S(X, p) := g^S(p, X)$.

For $X, Y \subseteq S$, the geodesic distance between X and Y in S , denoted $g^S(X, Y)$, is defined as

$$g^S(X, Y) := \bigwedge_{p \in Y} g^S(p, X)$$

Remark 10.3. Because geodesic distance is defined as an infimum over one-dimensional Hausdorff measures, it is determined by the lengths of paths or curves. Higher-dimensional sets have infinite H^1 measure and so are irrelevant to the value of the infimum.

The remainder of the definition of the geodesic distance function specifies its value when one or more of its arguments is a set of point, rather than a point. As with the Euclidean distance function, the geodesic distance between a point and a set or between two sets is an infimum of geodesic distances between pairs of points.

We now define a metric in the interest of characterizing the geodesic distance function as a metric:

Definition 10.2. A function $\Psi : X \times X \rightarrow \mathbb{R}^{\geq}$ is a *metric* if and only if for all $x, y, z \in X$

- (1) $\Psi(x, y) = 0 \Leftrightarrow x = y$
- (2) $\Psi(x, y) = \Psi(y, x)$
- (3) $\Psi(x, z) \leq \Psi(x, y) + \Psi(y, z)$

Claim 10.1. For path-connected $S \subseteq \mathbb{R}^d$, the geodesic distance function $g^S : S \times S \rightarrow \mathbb{R}^{\geq}$ is a metric. We consider verification of the first two conditions to be straightforward. For the third, we note that the set of path-connected sets that each contain x and z is a superset of the path-connected sets that each contain x, y , and z , so the length infimum over paths connecting x to z is necessarily less than or equal to the infimum over the smaller set of paths that also pass through y .

FIGURE 10.1. Lorem ipsum dolor sit amet

Remark 10.4. Having established that the geodesic distance function for a path-connected set is a metric, we define geodesic balls and the topology induced by the geodesic metric.

Notation 7. For $S \subseteq \mathbb{R}^d$, let $g_r^S(p)$ denote the open geodesic ball in S of radius $r \in \mathbb{R}^>$ centered at $p \in S$:

$$g_r^S(p) := \{q \in S \mid g^S(p, q) < r\}$$

Remark 10.5. Geodesic balls behave in many ways like their Euclidean counterparts. This is particularly true when the radius of a ball is small, which is typical in our work. When the radius of the geodesic neighborhood of a point on a surface is large, the neighborhood may “join itself,” and so have topology other than that of an open disk, such as the topology of an open annulus. In the plane, the singularities of the Euclidean distance function to a finite set of point sites partition the plane into a finite set of Voronoi regions, each consisting of points closest to a particular site. In contrast, the geodesic distance function on a smooth surface in \mathbb{R}^3 may have a singularity such that points on either side of a singularity are closest to the same site (Figure 10.1.)

Notation 8. Let the $\mathbb{G}\mathbb{T}$ denote the set of all geodesic balls in $T \subseteq \mathbb{R}^d$, and let the notation for specifying sets of balls in Euclidean described in Notation 2 apply to sets of geodesic balls. For instance, let $\mathbb{G}\mathbb{T}^S$ denote the set of geodesic balls in S .

Definition 10.3. For a path-connected set $T \subseteq \mathbb{R}^d$ and a set $S \subseteq T$, the *geodesic interior* of S in T , denoted $S \circ^T$, is defined as

$$S \circ^T := \bigcup \mathbb{G}\mathbb{T}^S$$

The *geodesic complement* of S in T is defined as $T \setminus S$.

The *geodesic closure* of S in T , denoted $S \bullet^T$, is defined as

$$S \bullet^T := T \setminus ((T \setminus S) \circ^T)$$

The *geodesic boundary* of S in T , denoted $\partial^T S$, is defined as

$$\partial^T S := S \bullet^T \cap ((T \setminus S) \bullet^T)$$

Remark 10.6. Given that T is path connected, the geodesic closure operator defines a closure structure. More specifically, the geodesic distance function for T induces a metric topology on T . Statements combining the geodesic topology on a set with Euclidean topology can be counterintuitive. For instance, a finite-radius open geodesic ball centered at a point on a line is an open line segment — a line segment that does not include its endpoints. However, it contains no open Euclidean ball, so its Euclidean interior is empty. Because it is not equal to its interior, the open line segment is not an open set in the Euclidean topology.

Remark 10.7. Although it is possible to define a geodesic morphology or geodesic medial axis, geodesic dilation is the operation most relevant to our contributions, because it lets us express integrals over closed sets as limits of integrals over their dilations when the dilation radius goes to zero.

Definition 10.4. The *geodesic dilation* of S in T by radius $r \in \mathbb{R}^>$, denoted $S \uparrow_r^T$, is defined as

$$S \uparrow_r^T := \bigcup \mathbb{G}\mathbb{T}_r(S) = \{p \in T \mid g^T(p, S) < r\}$$

10.2.3. Hausdorff dimension and local dimensionality.

Definition 10.5. For $S \subseteq \mathbb{R}^d$, the *Hausdorff dimension* of S , denoted HD^S , is defined as

$$\text{HD}^S := \bigwedge_{d \in \mathbb{R}^{\geq}} \text{H}^d(S) = 0$$

Remark 10.8. Note that while we primarily work with integer dimensions, the infimum in the definition of Hausdorff dimension is taken over reals. For a curve such as a circle C in three dimensions, $\text{H}^d(C) = 0$ over $d \in (1, \infty)$. Because we take the infimum of $(1, \infty)$, we obtain the expected dimensionality $\text{HD}^C = 1$.

Definition 10.6. The *local dimensionality* of $p \in \mathbb{R}^d$ in $S \subseteq \mathbb{R}^d$, denoted LD_p^S , is defined as

$$\text{LD}_p^S := \bigwedge_{d \in \mathbb{R}^{\geq}} \exists r \in \mathbb{R}^> : \forall \rho \in (0, r) : \text{H}^d(g_\rho^S(p)) = 0$$

Remark 10.9. We introduce local dimensionality to facilitate the description of dimensionally heterogeneous sets — a solid ball with a protruding line segment, for instance. The local dimensionality at a point on a set is equal to that of the highest-dimensionality part of its neighborhood. At a point inside the ball, the local dimensionality is three; at a point in the interior of the line segment, the local dimensionality is one; while at the point where the line segment intersects the ball the local dimensionality is three because the geodesic neighborhood of the intersection point includes a portion of the solid ball.

10.3. Normals for arbitrary subsets of Euclidean space

10.3.1. Overview. We build our definition of the normals to a subset of Euclidean space in three steps. First, we define normals to morphologically closed sets. Second, we incorporate a limiting process into this definition to obtain normals to topologically closed sets. Third, we take the union of normals to connected components of S and S^c that lie in a ball-shaped neighborhood of p to obtain normal fields for arbitrary subsets of Euclidean space. These fields are meant to be meaningful or

reasonable under diverse circumstances. Our general normal field definition, for instance, provides normals orthogonal to a line segment subtracted from a disk (Figure .) It also assigns a sphere of normals to each point in a point cloud, so that measures of normal variation over a dense cloud of points are infinite.

Measures of normal variation like integral Gaussian curvature only behave consistently across smooth and nonsmooth sets if we allow a normal at a point to be “used” more than once. Instead of a set of normals at a point, we define a nonnegative normal count for each direction. The normal count is greater than one only at points that are singular in two senses: they are both nonsmooth and possess a geodesic neighborhood whose boundary has more than four critical points. We conclude by defining the total number of times a normal is used over a region.

10.3.2. Normals to morphologically closed sets.

Remark 10.10. The definition below specifies whether or not a vector is an outward-pointing unit normal to a morphologically closed set $S \subseteq \mathbb{R}^d$ at a given point in space. If S is topologically regular, at point $p \in \partial S$ where there exists a unique tangent plane to ∂S at p , we assign a unique outward-pointing normal to S at p . Under our definition, the dimensionality of the set of normals to S is greater at points where S the tangent plane is not unique. If $S \subseteq \mathbb{R}^2$ is two-dimensional, there is a fan of normals at a convex corner of S , while if $S \subseteq \mathbb{R}^3$ there is a cone of normals at a convex corner and a fan of normals at each point along a convex edge. There are no normals to S at points in its interior and exterior.

Our definition implicitly associates a set of normals to S with each point in space. We provide a notation for this set of normals in Notation 10, after we complete our definition of normals to arbitrary subsets of Euclidean space.

Definition 10.7. For $S \subseteq \mathbb{R}^d$ and $r \in \mathbb{R}^>$, suppose S is morphologically r -closed. Then $v \in \mathbb{S}^{d-1}$ is a *normal* to S at $p \in \mathbb{R}^d$ if and only if $p + rv \notin S \uparrow_r$.

Remark 10.11. Suppose a set $S \subseteq \mathbb{R}^d$ is morphologically r -closed. Its dilation $S \uparrow_r$ by r is topologically open. The offset of a point $p \in \partial S$ along a vector v of length r lies inside $S \uparrow_r$ unless v is normal to S . If v is normal to S , the point $p + v$ lies on the boundary of $S \uparrow_r$, rather than within it.

Claim 10.2. Suppose $S \subseteq \mathbb{R}^d$ is both r -closed and topologically regular. Also suppose that point $p \in \partial S$ has a unique tangent plane ϖ to ∂S at p , with unit vector $v \in \mathbb{S}^{d-1}$ orthogonal to ϖ such that for sufficiently small $\epsilon \in \mathbb{R}^>$, we have $p + \epsilon v \notin S \uparrow_r$. Then v is the unique vector such that $p + rv \notin S \uparrow_r$. If for some $w \in \mathbb{S}^{d-1}$ where $w \neq v$ we have $p + rw \notin S \uparrow_r$, then there exists an open ball of radius r and center $p + rw$ tangent to ∂S at p on a plane orthogonal to w . This contradicts our assumption that ϖ is the unique tangent plane to ∂S at p .

Suppose $S \subseteq \mathbb{R}^3$ is a convex polyhedron (the convex hull of a finite set $V \subseteq \mathbb{R}^3$ of points such that $\text{AH}(V) = \mathbb{R}^3$ or $\text{HD}(\text{CH}(S)) = 3$; see Definition 11.7) Let $p \in \partial S$ lie in the interior of an edge e of S . Then the set of normals at p subtends an arc on \mathbb{S}^2 of length $\pi - \alpha$, where α is the internal dihedral angle between the two faces of S incident on e . Every plane through p disjoint from S contains e . We can parameterize the set of these planes by an angle in $[0, \pi)$. Out of that interval, a set of the planes covering an interval of length α intersect S , so the planes through p disjoint from S can be parameterized by an angular interval of length $\pi - \alpha$. For any such plane ϖ and $r \in \mathbb{R}^>$, there exists an open ball b of radius r and with $p \in \partial b$ such that $b \cap S = \emptyset$. If b has center q , then $q \notin S \uparrow_r$. Because $\|\vec{pq}\| = r$, the vector $\frac{1}{r}\vec{pq}$ is a normal to S , because $p + r\left(\frac{1}{r}\vec{pq}\right) = q$. Consequently, the set of normals to S at p subtends an arc of length $\pi - \alpha$ on \mathbb{S}^2 .

A similar argument establishes that if $p \in \partial S \cap V$ is a vertex of S , then the set of normals forms a region of area $2\pi - \alpha$, where α is the internal solid angle at p . We parameterize the set of all planes through p by a set of directions $D \subseteq \mathbb{S}^2$ that covers half of the unit sphere, so that D has an area of 2π . A subset of those planes covering

FIGURE 10.2. Lorem ipsum dolor sit amet

a solid angle of α intersects S . As above, each of the planes through p disjoint from S defines a normal to S at p .

10.3.3. Normals to topologically closed sets.

Remark 10.12. If we apply our definition of the normals to a morphologically r -closed set to topologically closed sets by letting r approach zero, we can obtain well-behaved normals at many boundary points. We do not, however, obtain the fan of normals at a concave corner of a two-dimensional set $S \subseteq \mathbb{R}^2$ that should match the fan of normals we obtain at a convex corner of S (Figure 10.2.) To correct this, we in essence perform a morphological closure of our input set $S \subseteq \mathbb{R}^d$ with a small radius $\epsilon \in \mathbb{R}^>$ and associate with each point $p \in \mathbb{R}^d$ the set of all normals to $S \bullet_\epsilon$ within a small neighborhood of p . Because the morphological closure by ϵ fillets concavities, we obtain the desired normals at concavities.

The definition below uses two parameters: one controls the size of a open-ball neighborhood, while the other controls the radius used for morphological closing. Paraphrasing the definition, we define a vector $v \in \mathbb{S}^{d-1}$ to be a normal to a topologically closed set $S \subseteq \mathbb{R}^d$ at point $p \in \mathbb{R}^d$ if there is a ball-shaped neighborhood of p with radius r such that for every smaller neighborhood of p of radius $\rho < r$ there exists a radius for morphological closing δ such that for every smaller radius $\epsilon < \delta$, vector v is a normal to the morphologically ϵ -closed set S at some point within the neighborhood of radius ρ .

Qualitatively, the effect is similar to a two-step process:

- (1) We close S with a very small radius, which introduces new boundary patches at concave singularities.
- (2) We concentrate the normals on those patches at the singularities themselves.

Informally, we seek the limit of the normal field for $S \bullet_\epsilon$ as $\epsilon \rightarrow 0$. The definition below formalizes that idea.

FIGURE 10.3. Lorem ipsum dolor sit amet

Definition 10.8. Suppose $S \subseteq \mathbb{R}^d$ is topologically closed. Then $v \in \mathbb{S}^{d-1}$ is a *normal* to S at $p \in \mathbb{R}^d$ if and only if there exists an $r \in \mathbb{R}^>$ such that for all $\rho \in (0, r)$ such that there exists a $\delta < \rho$ such that for all $\epsilon \in (0, \delta)$,

$$\exists q \in B_\rho(p) : q + \epsilon v \notin S \bullet_\epsilon \uparrow_\epsilon$$

Remark 10.13. Morphological closure of S by ϵ may introduce into ∂S patch from a sphere $\partial B_\epsilon(p)$. Note that

$$p \notin (\partial B_\epsilon(p)) \uparrow_\epsilon$$

Because the distance of p from $\partial B_\epsilon(p)$ is equal to ϵ , not less than or equal to ϵ . A vector scaled by ϵ that maps a point on $\partial B_\epsilon(p)$ to p is consequently a normal to $\partial B_\epsilon(p)$. Figure 10.3 illustrates this in the plane, where $\partial B_\epsilon(p)$ is a circle.

10.3.4. Normals to arbitrary subsets of Euclidean space.

Remark 10.14. When $S \subseteq \mathbb{R}^d$ is neither morphologically nor topologically closed, we define the normals to $S \in \mathbb{R}^d$ at $p \in \mathbb{R}^d$ in terms of connected components of S and its complement lying within a small ball-shaped neighborhood of p . Specifically, vector $v \in \mathbb{S}^{d-1}$ is a normal to S at p if there is a neighborhood $B_r(p)$ such that in all smaller neighborhoods, vector v is normal to the closure of a connected component S in that neighborhood, or $-v$ is normal to the closure of a connected component of S^c in the neighborhood.

It is possible for the closure of a connected component of S within a ball-shaped neighborhood to fill the entire neighborhood, despite the fact that the component is not itself a ball. A complete discussion of this situation is beyond the scope of this dissertation, but we can obtain this kind of behavior when S is “densely woven” or

“packed with strands,” as when

$$S = \bigcup \{\overline{pq} \mid p, q \in \mathbb{Q}^d \cap B_1(o)\}$$

For some origin $o \in \mathbb{R}^d$ with $d > 2$. Then the geodesic distance between two points in the network of strands closely approximates the Euclidean distance between them. We define every unit vector as normal to p under these conditions.

Definition 10.9. For $S \subseteq \mathbb{R}^d$ and $p \in \mathbb{R}^d$, a vector $v \in \mathbb{S}^{d-1}$ is *normal* to S at p if and only if there exists an $r \in \mathbb{R}^>$ such that for all $\rho \in (0, r)$

- (1) There exists an $x \in \kappa(S \cap B_\rho(p))$ such that v is normal to x^- at a point in $B_\rho(p)$.
- (2) There exists an $x \in \kappa(S^c \cap B_\rho(p))$ such that $-v$ is normal to x^- at a point in $B_\rho(p)$.
- (3) Neither $B_\rho(p) \subseteq S$ nor $B_\rho(p) \subseteq S^c$, but $(B_\rho(p) \cap S)^- = (B_\rho(p) \cap S^c)^- = B_\rho(p)^-$.

Theorem 10.1. For $S \subseteq \mathbb{R}^d$, a normal exists at $p \in \mathbb{R}^d$ if and only if $p \in \partial S$.

PROOF. Suppose $p \in \partial S$. Then neither $B_\rho(p) \subseteq S$ nor $B_\rho(p) \subseteq S^c$, so $\kappa(S \cap B_\rho(p))$ and $\kappa(S^c \cap B_\rho(p))$ are nonempty. If condition (1) or condition (2) from Definition 10.9 is true, there is a normal to S at p . Otherwise, then for every component $x \in \kappa(S \cap B_\rho(p)) \cup \kappa(S^c \cap B_\rho(p))$, we have $(x^c \cap B_\rho(p))^\circ = \emptyset$, because if $(x^c \cap B_\rho(p))^\circ \supseteq \emptyset$, there would be a closure normal defined at a point $q \in B_\rho(p) \cap x^-$, satisfying (1) or (2). If $(x^c \cap B_\rho(p))^\circ = \emptyset$, then $x^- = B_\rho(p)^-$. Because x is any connected component of $S \cap B_\rho(p)$ or $S^c \cap B_\rho(p)$, we conclude that condition (3) holds. We conclude that if $p \in \partial S$, there is a normal to S at p .

If $p \in S^\circ$, then there exists an $r \in \mathbb{R}^>$ such that for all $\rho \in (0, r)$, we have $S \cap B_\rho(p) = B_\rho(p)$. The normals to $B_\rho(p)$ lie on its boundary, and because $B_\rho(p)$ does not include its boundary, there is no normal to $B_\rho(p)$ satisfying condition (1.)

The set $S^c \cap B_\rho(p)$ is empty, so there are no normal satisfying condition (2.) Because $B_\rho(p) \subseteq S$, condition (3) is unsatisfied as well, and we conclude that there is no normal to S at $p \in S^\circ$. The proof that there is no normal to S at $p \in S^{oc}$ is essentially identical, leading us to conclude that there is no normal to S at p if $p \notin \partial S$. Combined with our proof that there is always a normal at p if $p \in \partial S$, we conclude a normal exists at $p \in \mathbb{R}^d$ if and only if $p \in \partial S$. \square

Remark 10.15. Now that we have defined normals to arbitrary subsets of Euclidean space, we introduce a notation for the set of normals to a set at a point or over a set of points. The set of normals to a topologically regular set $S \subseteq \mathbb{R}^d$ is empty in the interior and exterior of S , it contains a single vector where there is a unique tangent plane to ∂S , and it contains a fan or cone of normals at singularities such as cones or edges. At a point $p \in \partial S$, the set of normals to S at p may contain one vector or it may contain all vectors in \mathbb{S}^{d-1} .

We define a set's normal field as the set containing all of its point-normal pairs. We use subsets of the normal field to specify maximal path-connected components of point-normal pairs, over which both position and normal vary continuously. Our motivation is that, under our normal field definitions, the boundary of a topologically regular solid has only outward-pointing normals, but its boundary has both inward- and outward-pointing normals. A ball, for instance, has one normal at every boundary point, and that normal points away from the center of the ball. The ball's boundary, however, has two normals at each point, one pointing toward the ball's center and the other pointing away. This can lead to what we consider undesirable consequences for our measures of normal variation, such as that the total mean curvature of a sphere using its entire normal field is zero, but the total mean curvature for a three-dimensional ball of radius r is $4\pi r$. The total mean curvature of the ball is the total mean curvature of the sphere in classic differential geometry: it is the product of the mean curvature $\frac{1}{r}$ and the sphere's surface area of $4\pi r^2$.

By providing means to refer to path-connected sets of normals, we enable the extraction of the normals from either side of an oriented surface. For instance, given an outward-pointing normal v at a point p on a sphere, the point-normal pairs that are path-connected to (p, v) only include outward-pointing normals. Under our definitions, taking the the mean curvature over a sphere using only its outward-pointing normals yields the classic value of $4\pi r$.

Notation 9. For $S \subseteq \mathbb{R}^d$, let N_p^S denote the set of all $v \in \mathbb{S}^{d-1}$ such that v is a normal to S at p .

Definition 10.10. For $S \subseteq \mathbb{R}^d$, the the *normal field* to S , denoted NF^S is the set of all point-normal pairs

$$NF^S := \{(p, v) \mid (p \in \partial S) \wedge (v \in N_p^S)\}$$

Notation 10. For $S \subseteq \mathbb{R}^d$, $X \subseteq \partial S$, and $Y \subseteq NF^S$ let $N_{X,Y}^S$ denote the set of all $v \in \mathbb{S}^{d-1}$ such that v is a normal to S at some point $p \in X$ such that (p, v) is path-connected to a point-normal pair $x \in Y$.

Let $N_X^S := N_{X, NF^S}^S$, let $N_{p,Y}^S := N_{\{p\}, Y}^S$, and let $N_{X, (p,v)}^S := N_{X, \{(p,v)\}}^S$.

Remark 10.16. If $S \subseteq \mathbb{R}^3$ is an orientable $(d - 1)$ -manifold, then to obtain either the outward- or inward-pointing subset of the normal field of S requires a single point-normal pair. With a set of point-normal pairs, we can specify the union of inward- and outward-pointing normal fields for a single connected component of S , as well as combinations of inward- and outward- pointing fields from different components of S .

10.3.5. Smooth and singular points.

Definition 10.11. For $S \subseteq \mathbb{R}^d$, a point $p \in \partial S$ is *smooth* if and only if the sum of the local dimensionality of ∂S and the dimensionality of N_p^S is equal to $d - 1$. Set S is smooth if and only if for all $p \in \partial S$ the set ∂S is smooth.

Point $p \in S$ is *singular* if and only if it is not smooth.

FIGURE 10.4. On the left, we show a star-shaped set. Treating the set as lying in a horizontal plane, on the right we show the result of tapering it as we sweep it vertically, generating a solid that might be described as a pleated cone.

Remark 10.17. Our definition of smoothness is less restrictive than alternative in use that require that the curvature of the boundary of a surface be continuous or infinitely differentiable. By contrast, we define a surface in three dimensions as smooth if and only if it has a unique tangent plane at every point. Our definition is also applicable to subsets of \mathbb{R}^d with dimensionality less than $d - 1$. For instance, if S is a planar circle in \mathbb{R}^3 , then N_p^S is a great circle in \mathbb{S}^2 .

10.3.6. Normal count.

10.3.6.1. Normal count at a point.

Remark 10.18. Given $S \subseteq \mathbb{R}^d$ and $p \in \partial S$, measures like the slack at p (Definition 10.14) or the Gaussian curvature at p (Definition 10.17) measure the “quantity” of normals in the neighborhood of p . If we take each normal in the neighborhood of p to be a point on the sphere \mathbb{S}^{d-1} , slack and Gaussian curvature measure the mapping of the normals in the neighborhood of p onto the sphere. If we measure slack or Gaussian curvature by taking the measure of the points on the sphere corresponding to the set of normals to S at p , we generate inconsistencies involving the maximum amount of normal variation that can occur in a point’s neighborhood. The set of normals to S at p can potentially contain every vector in \mathbb{S}^{d-1} , in which case the measure of the points corresponding to those normals over a neighborhood of p would equal the measure of \mathbb{S}^{d-1} — 2π for the circle in two dimensions and 4π for the sphere in three dimensions. The problem is that slack and absolute Gaussian curvature can be arbitrarily high in the neighborhood of a point on a set’s boundary, because the mapping from the normals in the point’s neighborhood to the sphere can “fold over,” “covering” points on the sphere multiple times (Figure 10.4.)

To preserve topological invariants like the Gauss-Bonnet Theorem, we assign each normal a nonnegative weight we call its normal count, which corresponds to the number of times that normal is “used” or in a given neighborhood. We then compute slack and Gaussian curvature using areas on the sphere that are weighted by normal count.

We note that a normal count greater than one is singular in the sense that it can be eliminated by a small perturbation. The existence of normal counts greater than one in solid models, however, is not implausible. The set shown in Figure 10.4 might appear in a model of a gear or citrus juicer.

We define the normal count of a vector $v \in \mathbb{S}^{d-1}$ to $S \subseteq \mathbb{R}^d$ at $p \in \mathbb{R}^d$ by slightly dilating the connected components of S an S^c in a small neighborhood of p . For each dilated component k of S , we count the number of connected components of points from k that have v as one of their normals. We take the sum of that count over all the dilated components, then compute an analogous number using S^c . The normal count of v is the least upper bound on those two numbers as the neighborhood size and dilation radius vanish. The small dilation functions as a perturbation that “breaks up” the multiple normal uses at a vertex into separates them into patches that each use the normal once. We count the patches to obtain the normal count.

Definition 10.12. For $S \subseteq \mathbb{R}^d$, the *normal count* of vector $v \in \mathbb{S}^{d-1}$ to S at $p \in \mathbb{R}^d$, denoted $\text{NC}_p^S(v)$, is a cardinal number. If neither $B_\rho(p) \subseteq S$ nor $B_\rho(p) \subseteq S^c$, but $(B_\rho(p) \cap S)^- = (B_\rho(p) \cap S^c)^- = B_\rho(p)^-$, we define $\text{NC}_p^S(v) := \aleph_0$. Otherwise, $\text{NC}_p^S(v)$ is the least cardinal over all $r \in \mathbb{R}^>$ such that for all $\rho \in (0, r)$ there is a $\delta < \rho$ such that for all $\epsilon \in (0, \delta)$

- (1) For $X = \kappa(S \cap B_\rho(p))$, we have $\text{NC}_p^S(v) \geq \sum_{k \in X} \#\kappa(\{p \in \mathbb{R}^d \mid v \in N_p^{k \uparrow \epsilon}\})$
- (2) For $Y = \kappa(S^c \cap B_\rho(p))$, we have $\text{NC}_p^S(v) \geq \sum_{k \in Y} \#\kappa(\{p \in \mathbb{R}^d \mid -v \in N_p^{k \uparrow \epsilon}\})$

10.3.6.2. Normal count over a region.

FIGURE 10.5. Suppose S is a cube and F is a face of S . The set F is a connected component of points that share the same normal v . Accordingly, we refer to F as a normal patch of S . If $X \subseteq \partial S$ has a “frilly” boundary, the number of connected components in the intersection between X and F may be arbitrarily high, so counting those connected components can overestimate normal variation. Instead, we count over the patches that X intersects, so F counts as a single use of vector v .

Remark 10.19. The number of uses of a normal $v \in \mathbb{S}^{d-1}$ to $S \subseteq \mathbb{R}^d$ over a region $X \subseteq \partial S$ is not equal to the sum of the normal count $\text{NC}_p^S(v)$ over all $p \in X$. If it were, then the face of a cube would have an infinite number of uses of the same vector, despite the fact that the normal variation of the face — which is what we seek to measure — is zero. Instead, we decompose the normals to S into connected components that share the same normal. The normal count for a vector over the region X depends both on the number of patches with that normal that X intersects as well as the greatest normal count within the intersection of X and a patch. As illustrated in Figure , counting over the components of the intersection between X and the normal patches from S rather than over the patches that X intersects can overestimate normal variation. While our definition of normal count is novel, many of these concepts follow from Kuiper’s discussion of topsets [37].

Definition 10.13. Let the set of *normal patches* of $S \subseteq \mathbb{R}^d$ for vector $v \in \mathbb{S}^{d-1}$ from $Y \subseteq \text{NF}^S$, denoted $\text{NP}_Y^S(v)$, be defined as

$$\text{NP}_Y^S(v) := \kappa(\{p \in \mathbb{R}^d \mid v \in \text{N}_{p,Y}^S\})$$

Let the *normal count* of vector $v \in \mathbb{S}^{d-1}$ to S over $X \subseteq \partial S$ from $Y \subseteq \text{NF}^S$, denoted $\text{NC}_{X,Y}^S(v)$, be defined as

$$\text{NC}_{X,Y}^S(v) = \sum_{k \in \text{NP}_Y^S(v)} \left[\bigvee \{ \{ \text{NC}_p^S(v) \mid p \in k \cap X \} \cup \{0\} \} \right]$$

10.4. Slack and Gaussian curvature

10.4.1. Overview. We define slack and Gaussian curvature, both of which measure normal variation. Informally, the “normal variation” over a surface patch is the quantity of distinct normals associated with the patch. It is zero on a planar patch, where the normal is the same at every point, and it increases in a fixed-radius geodesic neighborhood of a point on a sphere as the radius of the sphere decreases.

We begin by defining slack, which is a generalized form of absolute Gaussian curvature. The slack over a boundary patch is the integral over the unit sphere of the normal counts for the image of that patch on the sphere. Our definition of Gaussian curvature is based on our definition of slack, but incorporates information about how the motion of a point’s neighborhood under its normals reverses or preserves its orientation. While slack is nonnegative, negative orientation values corresponding to neighborhood orientation reversals can yield negative Gaussian curvatures.

We close by considering Gaussian curvature on triangle meshes. We show that Gaussian curvature concentrates at mesh vertices, and it can be computed by summing the angles incident on a vertex. As a result, a mesh’s total slack and Gaussian curvature can be computed with a linear-time mesh traversal.

10.4.2. Slack.

Definition 10.14. For $S \subseteq \mathbb{R}^d$, $X \subseteq \partial S$, and $Y \subseteq \text{NF}^S$, the *slack* of S over X from Y , denoted $\text{SL}_{X,Y}^S$, is defined as

$$\text{SL}_{X,Y}^S := \sum_{i=1}^{\infty} \text{H}^{d-1}(\{v \in \mathbb{S}^{d-1} \mid \text{NC}_{X,Y}^S(v) \geq i\})$$

The slack of S , denoted SL^S , is defined as $\text{SL}^S := \text{SL}_{\partial S, \text{NF}^S}^S$.

The slack of S at $p \in \partial S$ from $X \subseteq \text{NF}^S$, denoted $\text{SL}_{p,X}^S$, is defined as

$$\text{SL}_{p,X}^S := \lim_{r \rightarrow 0} \frac{\text{SL}_{g_r^{\partial S}(p), X}^S}{\text{H}^{d-1}(g_r^{\partial S}(p))}$$

Remark 10.20. Once we define the slack of a set over a subset of its boundary, the slack of the set itself is the slack over its boundary. The slack at a point is closely related to the absolute value of the Gaussian curvature at that point. Both measure the ratio between the area of a neighborhood of a point and the area of the neighborhood's image under the Gauss map as the neighborhood size goes to zero. Under the definition of Gaussian curvature we provide in Definition 10.17, slack and the absolute value of Gaussian curvature are identical at any point where the sign function (Definition 10.15) over a small neighborhood of the point is constantly equal to either positive or negative one.

Suppose $S \subseteq \mathbb{R}^d$ is a $(d - 1)$ -dimensional orientable manifold with differentiable normals. If we use classic definitions of Gaussian curvature, then the Gaussian curvature at a point $p \in S$ is equal the product of the principal curvatures at p . The slack at p from $p \times N_p^S$ is twice as large as the absolute value of that product, because the slack value will count two normals with opposite directions at every point in the neighborhood of p . If instead we select a vector $v \in N_p^S$, the slack at p from (p, v) and the absolute Gaussian curvature at p are equal. If S bounds a $(d - 1)$ -dimensional set T , then the slack of T at p and the absolute Gaussian curvature of S at p are also equal. Although there are two normals to ∂T at each of its points, there is only one normal to T at each of its boundary points (Figure .)

10.4.3. Gaussian curvature.

Remark 10.21. While slack measures the area of the image of the neighborhood of a point under the Gauss map, Gaussian curvature measures signed area. The sign is obtained from orientation information. In the following definition, we consider a small geodesic neighborhood of a point and the normals over that neighborhood continuously reachable from a reference normal. We seek to determine whether the geodesic neighborhood “flips” if we advance neighborhood points sufficiently far along their normals. If the neighborhood flips, the sign is negative. Otherwise, it is positive.

We determine if the neighborhood flips by shooting a ray from each point in the neighborhood of p along each of its normals until it intersects a large sphere centered on p . The ray origins and sphere intersection points define a set of line segments connecting every point in the neighborhood p to a matching point from a patch on the sphere. Suppose we isotopically deform the patch on the sphere while maintaining the correspondences between the points in the spherical patch and the points in the neighborhood of p . We assert that if there exists a deformation such that the line segments “untangle” and no two segments intersect, then the neighborhood is not flipped. If the neighborhood is flipped, by contrast, the line segments cross regardless of how we deform the spherical patch.

Definition 10.15. For $S \subseteq \mathbb{R}^d$, point $p \in \partial S$, and vector $v \in N_p^S$, the sign of S at p for v , denoted $\sigma_p^S(v)$, is defined if and only if there exists a $\delta \in \mathbb{R}^>$ such that for all $\epsilon \in (0, \delta)$, the ϵ -geodesic neighborhood of p has an empty Euclidean interior: $g_\epsilon^S(p)^\circ = \emptyset$. If $\sigma_p^S(v)$ is defined, suppose that for all $\epsilon \in (0, \delta)$, there is an $r \in \mathbb{R}^>$ such that for all $\rho \in (r, \infty)$, there exists an isotopy $H : [0, 1] \times \partial B_\rho(p) \rightarrow \partial B_\rho(p)$ such that for all pairs $P \in \mathbb{R}^d \times \mathbb{R}^d$ given by

$$P := \left\{ (a, b) \mid a \in g_\epsilon^S(p) \wedge b \in \left\{ \overrightarrow{at} \cap \partial B_\rho(p) \mid t \in N_a^S \right\} \right\}$$

We have

$$\forall x, y \in \left\{ \overline{aH_1(b)} \setminus a \setminus H_1(b) \mid (a, b) \in P \right\} : x \cap y = \emptyset$$

Then $\sigma_p^S(v) := 1$. Otherwise, $\sigma_p^S(v) := -1$.

Example 10.1. Take S to be the unit-radius sphere: $S := \partial B_1(o)$, for some choice of origin $o \in \mathbb{R}^d$. At any point $p \in S$, the normal set at p consists of two vectors, one inward-pointing and one outward-pointing: $N_p^S := \{o - p, p - o\}$. For the outward-pointing vector, all sectional curvatures are one, while for the inward-pointing vector all sectional curvatures are negative one. The Gaussian curvature should be one

for either vector, so the sign for either vector should be one as well: $\sigma_p^S(o - p) = \sigma_p^S(p - o) = 1$. For vector $p - o$, which is outward-pointing, the union of the outward-pointing rays over a small geodesic neighborhood of p is set difference between a cone with apex at o and the unit ball $B_1(o)$. The ray at neighborhood point q has direction $o - q$ and is a subset of the ray from o with the same direction. As these rays are distinct for all neighborhoods on $\partial B_1(o)$, the line segments obtained by intersecting them with a sufficiently large ball centered on p are disjoint, leading to a positive sign: $\sigma_p^S(p - o) = 1$.

For the inward-pointing vector $o - p$, the union of the rays over a small geodesic neighborhood of p define a double-napped cone, with its singular point at o . The intersection of these rays with a large sphere $\partial B_\rho(p)$ centered on p is a geodesic disk G , but the segments connecting points on S with their corresponding ray intersection points all converge at o . We must ask whether there is an isotopy of G on $\partial B_\rho(p)$ that eliminates the segment intersections. Rotating G by about its center by π shifts the connecting segments into the same conical configuration seen in our explanation of $\sigma_p^S(p - o)$. As the segments no longer intersect, we have $\sigma_p^S(o - p) = 1$.

Example 10.2. Let $S \subseteq \mathbb{R}^3$ be the epigraph of the function $f : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ given by $f(x, y) = xy$. Let p be the origin, so $p = (0, 0, 0)$, and let v be the vector pointing in z -direction, so $v = (0, 0, 1)$. The rays over a geodesic neighborhood of the origin intersect along the plane ϖ orthogonal to the xy -plane that intersects the xy -plane at the line $y = -x$. For sufficiently large ρ , the ray segment from every point in the neighborhood of p crosses ϖ before intersecting $\partial B_\rho(p)$. To eliminate the intersections between segments, we would need to reflect the points on $\partial B_\rho(p)$; for instance, we could reflect them across ϖ . This is equivalent to asserting that the shape operator (Definition 11.2) has a negative determinant. Unlike the rotation used in Example 10.1, a reflection cannot be accomplished achieved by an isotopy, so $\sigma_p^S(v) = -1$.

The section of f through ϖ is a function $g(t) = -\frac{t^2}{2}$ with $t = \sqrt{x^2 + y^2}$, which has a curvature of -1 at p . Similarly, the section of f orthogonal to ϖ and the xy -plane

is $h(t) = \frac{t^2}{2}$, yielding a sectional curvature of 1. The Gaussian curvature of S at p from v is then -1 , agreeing with our sign of $\sigma_p^S(v)$.

Remark 10.22. The identification of point-normal pairs with a sign motivates the decomposition of a set's boundary into patches that share the same normal direction and sign. This decomposition provides us with positive and negative normal counts for each direction on the sphere. We interpret these spherical connected components of positive and negative counts as signed "normal areas."

Definition 10.16. For $S \subseteq \mathbb{R}^d$, the set of *positive normal patches* for vector $v \in \mathbb{S}^{d-1}$ from $X \subseteq \text{NF}^S$, denoted $\text{NP}_X^{S+}(v)$, is defined as

$$\text{NP}_X^{S+}(v) := \kappa \left(\{p \in \mathbb{R}^d \mid (v \in \text{N}_{p,X}^S) \wedge (\sigma_p^S(v) = 1)\} \right)$$

The set of *negative normal patches* for v from X , denoted $\text{NP}_X^{S-}(v)$, is defined as

$$\text{NP}_X^{S-}(v) := \kappa \left(\{p \in \mathbb{R}^d \mid (v \in \text{N}_{p,X}^S) \wedge (\sigma_p^S(v) = -1)\} \right)$$

For $X \subseteq \partial S$ and $Y \subseteq \text{NF}^S$, assume that for all $p \in X$ and vector $v \in \text{N}_{p,Y}^S$, the normal count $\text{NC}_p^S(v)$ is finite. Then the *signed normal count* of v to S over X from Y , denoted $\text{SC}_{X,Y}^S(v)$, is defined as

$$\begin{aligned} \text{SC}_{X,Y}^S(v) := & \sum_{k \in \text{NP}_Y^{S+}(v)} \left[\bigvee \{ \{ \text{NC}_p^S(v) \mid p \in k \cap X \} \cup \{0\} \} \right] \\ & - \sum_{k \in \text{NP}_Y^{S-}(v)} \left[\bigvee \{ \{ \text{NC}_p^S(v) \mid p \in k \cap X \} \cup \{0\} \} \right] \end{aligned}$$

Remark 10.23. The definitions of normal count and signed normal count have similar structures. In the definition of signed normal count, we decompose normal patches into those where the sign function is positive and those where it is negative. The signed count for a vector over a region is then the difference between the count from the positive patches and the count from the negative patches. The signed normal

count maps each point on the sphere to an integer, which may be positive, zero, or negative.

We obtain integral Gaussian curvature by integrating positive and negative normal counts separately over the sphere, and then taking their difference:

Definition 10.17. For $S \subseteq \mathbb{R}^d$, $X \subseteq \partial S$, and $Y \subseteq \text{NF}^S$, the *Gaussian curvature* of S over X from Y , denoted $G_{X,Y}^S$, is defined as

$$G_{X,Y}^S := \sum_{i=1}^{\infty} \text{H}^{d-1}(\{v \in \mathbb{S}^{d-1} \mid \text{SC}_{X,Y}^S(v) \geq i\}) \\ - \sum_{i=1}^{\infty} \text{H}^{d-1}(\{v \in \mathbb{S}^{d-1} \mid \text{SC}_{X,Y}^S(v) \leq -i\})$$

The *Gaussian curvature* of S , denoted $G^S := G_{\partial S, \text{NF}^S}^S$.

The *Gaussian curvature* of S at $p \in \partial S$ for vector $v \in \mathbb{S}^{d-1}$ from $X \subseteq \text{NF}^S$, denoted $G_{(p,v)}^S$, is defined as:

$$G_{(p,v)}^S := \lim_{r \rightarrow 0} \frac{G_{g_r^{\partial S}(p), X}^S}{\text{H}^{d-1}(g_r^{\partial S}(p))}$$

Remark 10.24. If $S \subseteq \mathbb{R}^3$ is a polyhedron, its Gaussian curvature is concentrated at its vertices, in the sense that the Gaussian curvature of S is zero over any set that does not contain a vertex of S . While the Gaussian curvature at a vertex of S is infinite, the Gaussian curvature over a small neighborhood of a vertex is finite, and the Gaussian curvature of S is equal to the sum over all vertices of S of the Gaussian curvature over the neighborhood of each vertex.

The Gaussian curvature over the neighborhood of a vertex is given by the angle defect formula. Suppose p is a vertex of S with n edges incident on it. If we order the edges cyclically, we obtain a sequence of angles $\alpha_1, \alpha_2, \dots, \alpha_n$ between consecutive edges, with each angle in the interval $(0, \pi]$. The angle defect formula states:

$$\lim_{r \rightarrow 0} G_{g_r^S}^S(p) = 2\pi - \sum_{i=1}^n \alpha_i$$

If S is represented on a computer as a triangle mesh, computing the value of the right side of the equation requires traversing the triangles incident on the vertex. Computing the sum of all such values, which is the mesh's total Gaussian curvature, requires a single traversal of the mesh.

10.5. Mean and sectional curvature

Remark 10.25. The following definition of mean curvature is based on the idea that “mean curvature is the gradient of surface area,” in the sense that if we “move the surface along” the field of mean curvature normals produces by the Laplace-Beltrami operator, its area decreases as quickly as possible out of all nearby surface perturbations.

Definition 10.18. For $S \subseteq \mathbb{R}^d$, $X \subseteq \partial S$, and $Y \subseteq \text{NF}^S$ and radius $r \in \mathbb{R}^>$, the *trimmed offset* of X for S from Y by r , denoted $X \uparrow_r^{S,Y}$, is defined as

$$X \uparrow_r^{S,Y} := \bigcup \left\{ \overline{p(p+rv)} \mid (p \in X) \wedge (v \in N_{p,Y}^S) \right\} \setminus X \uparrow_r$$

Remark 10.26. If $S \subseteq \mathbb{R}^d$ is a topologically regular set with differentiable normals and $X \subseteq \partial S$ is a patch from ∂S , then $X \uparrow_r^{S,\text{NF}^S}$ is the trimmed outward offset of X along the normals from S . If instead $S \subseteq \mathbb{R}^d$ is an orientable $(d-1)$ -dimensional manifold and X is a patch from ∂S , then each point of X has two normals, and it is displaced along both of them. Consequently, $X \uparrow_r^{S,\text{NF}^S}$ is the union of two offsets. Either may be designated as the inward offset, while the other would be the outward offset.

Definition 10.19. For $S \subseteq \mathbb{R}^d$, suppose $(\partial S)^\circ = \emptyset$, $X \subseteq \partial S$, and $Y \subseteq \text{NF}^S$. Then the *mean curvature* of S over X from Y , denoted $M_{X,Y}^S$, is defined as

$$M_{X,Y}^S := \lim_{\epsilon \rightarrow 0} \frac{H^{d-1}(X \uparrow_{\epsilon}^{S,Y}) - H^{d-1}(X \uparrow_{-\epsilon}^{S,Y})}{2\epsilon(d-1)}$$

The *mean curvature* of S , denoted M^S , is defined as $M^S := M_{\partial S, \text{NF}^S}^S$.

The *mean curvature* of S at $p \in \partial S$ from $X \subseteq \text{NF}^S$, denoted $M_{p,X}^S$, is defined as

$$M_{p,X}^S := \lim_{r \rightarrow 0} \frac{M_{g_r^{\partial S}(p), X}^S}{H^{d-1}(g_r^{\partial S}(p))}$$

Remark 10.27. If $S \subseteq \mathbb{R}^3$ is a solid polyhedron, its mean curvature is concentrated on its edges, in the sense that the integral mean curvature of S is zero over any set that does not include an edge of S . Note that such a set would consist of one or more planar patches from the faces of S . As with Gaussian curvature, the mean curvature at a point on an edge is infinite, but the integral mean curvature over a small neighborhood of an edge is finite, and the total mean curvature of S is equal to the sum over the edges of S of the integral mean curvature over the neighborhood of each edge.

The integral mean curvature over a small neighborhood of an edge \overline{pq} is proportional to the edge's length and the angle α formed by the two faces incident on the edge:

$$\lim_{r \rightarrow 0} M_{g_r^{\partial S}(\overline{pq})}^S = \frac{(\pi - \alpha) \|p - q\|}{2}$$

The factor of 2 in the formula's denominator is due to the fact that the edge bends along only one direction. Recall that the mean curvature of a smooth surface in three dimensions is the average of its maximum and minimum sectional curvatures, whose corresponding sections are orthogonal. At a mesh edge, the maximum sectional curvature is orthogonal to the edge, and its value is determined by the change in normal direction as we move across the edge. That value is precisely the angle between the faces incident on the edge. The minimum sectional curvature, by contrast, lies along the edge. The section of the edge is a line, so the sectional curvature through the

edge is zero. The mean curvature over a neighborhood of the edge is then proportional to half the change in angle across the edge.

If S is the boundary of a polyhedron rather than a polyhedron, then the mean curvature over the neighborhood of an edge is zero when we compute it from the entire normal field NF^S . The two sets $X \uparrow_{\epsilon}^{S, \text{NF}^S}$ and $X \uparrow_{-\epsilon}^{S, \text{NF}^S}$ are identical, so the difference between their Hausdorff measures taken in Definition 10.19 is zero. To get a more meaningful mean curvature value, we must designate a subset of the edge neighborhood's as outward-pointing. For instance, X is our edge neighborhood and we identify (p, v) as a point-normal pair with outward-pointing normal v , we compute $M_{X, (p, v)}^S$ rather than M_{X, NF^S}^S .

We now formalize the concept of sectional curvature:

Definition 10.20. For $S \subseteq \mathbb{R}^d$ and $p \in \partial S$ with $v \in N_p^S$, let ϖ be a two-dimensional affine subspace containing the line $\overleftrightarrow{p(p+v)}$. Also let $\Psi : \mathbb{R}^d \rightarrow \mathbb{R}^2$ isometrically map ϖ to \mathbb{R}^2 . Then the *sectional curvature* of S at p from v through ϖ , denoted $\text{SC}_p^S(\varpi)$, is defined as

$$\text{SC}_{p,v}^S(\varpi) := M_{\Psi(p), (\Psi(p), \Psi(v))}^{\Psi(S \cap \varpi)}$$

The *sectional curvature interval* for S at p from v , denoted $\text{SCI}_{p,v}^S$ is the set of sectional curvatures over all sections of S at p from v :

$$\text{SCI}_{p,v}^S := \left\{ \text{SC}_{p,v}^S \left(\text{AH} \left(\overleftrightarrow{p(p+v)} \cup \{q\} \right) \right) \mid (v \in N_p^S) \wedge (q \notin \text{AH} \left(\overleftrightarrow{p(p+v)} \right)) \right\}$$

The *maximum curvature* of S at p from v is defined to be $\bigvee \text{SCI}_{p,v}^S$, while the *minimum curvature* of S at p from v is defined to be $\bigwedge \text{SCI}_{p,v}^S$.

Remark 10.28. On the boundary of a topologically regular set with differentiable normals, our one-dimensional definition of mean curvature coincides with the standard definition of curvature. If $S \subseteq \mathbb{R}^3$ is a topologically regular set with differentiable

normals, we expect the maximum and minimum curvatures to lie in orthogonal sections.

If $S \subseteq \mathbb{R}^d$ is a $(d - 1)$ -dimensional manifold with differentiable normals, N_p^S contains two normals, which yield sectional curvatures with opposite signs.

10.6. Lipschitz continuity and smooth offset interval

Remark 10.29. Lipschitz continuity lets us bound the rate of change of the normals over the boundary of a set more simply and directly than we might using curvature.

Definition 10.21. For $S \subseteq \mathbb{R}^d$, $X \subseteq \partial S$, and $Y \subseteq \text{NF}^S$, the normals to S are ω -Lipschitz continuous over X from Y for $\omega \in \mathbb{R}^>$ if and only if for all $p, q \in X$,

$$\omega \geq \frac{g^{\mathbb{S}^{d-1}}(N_{p,Y}^S, N_{q,Y}^S)}{g^S(p, q)}$$

Let Ω denote the set of all ω such that the normals to S are ω -Lipschitz continuous over X . Then the *Lipschitz continuity* of the normals to S over X is equal to the infimum $\bigwedge \omega$ of the values in ω .

Remark 10.30. A set with $\frac{1}{r}$ -Lipschitz continuous normals has all sectional curvatures bounded by $\frac{1}{r}$, making it qualitatively rounded or blobby. While morphologically r -regular sets (Definition 9.2) have $\frac{1}{r}$ -Lipschitz continuous normals, r -regular sets are also “thick” — their interiors and exteriors are invariant under morphological r -opening. The Lipschitz continuity of a set’s normals provides no guarantee concerning its thickness.

Remark 10.31. The smooth offset interval at a point p measures how far a surface patch containing p can be offset inward or outward before developing a singularity at the offset position of p . The smooth offset interval is a set of real numbers, potentially including both positive and negative numbers. Negative numbers correspond to inward offsets.

In its most general form, the smooth offset interval is defined for a set $S \subseteq \mathbb{R}^d$, a boundary patch $X \subseteq \partial S$, a point $p \in X$, and a set Y specifying a maximal connected component of the normal field of S . The interval contains all $r \in \mathbb{R}$ such that for any normal v to S at p where (p, v) is path connected to Y , the point $p + rv$ lies outside the dilation of X by the absolute value of r . The smooth offset interval over a region is the intersection of the smooth offset intervals at each of its points. The smooth offset interval of S itself is the range of radii over which it can be offset without developing singularities, which is closely related to its regularity and least feature size. The smooth offset interval at a point provides information regarding curvature, but behaves differently from curvature measures like the sectional curvature interval.

Definition 10.22. For $S \subseteq \mathbb{R}^d$, $X \subseteq \partial S$, $p \in X$, and $Y \subseteq \text{NF}^S$, the *smooth offset interval* of S at p for X from Y , denoted $\text{SOI}_{p,X}^S$, is defined as

$$\begin{aligned} \text{SOI}_{p,Y}^{S,X}(p) := & \bigvee \{r \in \mathbb{R}^> \mid \forall v \in N_{p,Y}^S : p + rv \notin X \uparrow_r\} \cup \{0\} \\ & \cup \bigvee \{r \in \mathbb{R}^< \mid \forall v \in N_{p,Y}^S : p + rv \notin X \uparrow_{-r}\} \end{aligned}$$

The smooth offset interval of S over X from Y , denoted $\text{SOI}_{X,Y}^S$, is defined as

$$\text{SOI}_{X,Y}^S := \bigcap_{p \in X} \text{SOI}_{p,Y}^{S,X}$$

The smooth offset interval of S , denoted SOI^S , is defined as $\text{SOI}^S := \text{SOI}_{\partial S, \text{NF}^S}^S$.

The smooth offset interval of S at p from Y , denoted $\text{SOI}_{p,Y}^S$, is defined as

$$\text{SOI}_{p,Y}^S := \lim_{r \rightarrow \infty} \text{SOI}_{g_r^S(p), Y}^S$$

Remark 10.32. Suppose $S \subseteq \mathbb{R}^d$ is a topologically regular set with differentiable normals, the sectional curvature interval and smooth offset interval at $p \in \partial S$ provide identical information when the sectional curvature interval contains both positive and negative values. Then the maximum value in the smooth offset interval is the

negative of the reciprocal of the minimum value in sectional curvature interval, and the minimum value in the smooth offset interval is similarly the negative of the reciprocal of the maximum value in the sectional curvature interval. Suppose, however, that $S \subseteq \mathbb{R}^3$ is an ellipsoid. Then the minimum and maximum values of the sectional curvature interval are finite and positive. The maximum value for the smooth offset interval is instead unbounded, and the minimum smooth offset interval value is as before the negative of the reciprocal of the maximum sectional curvature.

Now suppose $S \subseteq \mathbb{R}^2$ has a boundary point $p \in \partial S$ where two circular arcs of equal radius r but opposite signed curvatures join with normal continuity. The sectional curvature interval at p is $[0]$, because that is the value of the mean curvature at p . The smooth offset interval, by contrast is $[-r, r]$. More generally, the smooth offset interval of a curved surface will contain a finite value. The sectional curvature interval for a curved surface may instead be $[0]$, as illustrated by the following example:

Example 10.3. Suppose \mathbb{R}^3 has a coordinate system with axes x , y , and z , and origin o . In the xy -plane, let C be the union of two arcs, each subtending $\frac{\pi}{2}$ radians. We take one arc from the circle of unit radius centered on $(0, 1)$, with the arc endpoints at $(-1, 1)$ and $(0, 0)$. Similarly, we take the other arc from the circle of unit radius centered on $(0, -1)$, with arc endpoints at $(1, -1)$ and $(0, 0)$. If we let $S \subseteq \mathbb{R}^2$ be the epigraph of C , then the mean curvature of S at $(0, 0)$ is zero.

We now rotate C about the y -axis while scaling it, sweeping a surface $X \subseteq \mathbb{R}^3$. In a plane through the y -axis forming an unsigned dihedral angle α with the xy -plane, we scale the y -coordinates of the points in C by $\frac{\pi}{2} - \alpha$. In the yz -plane, the scaled curve is then a line segment from $(0, -1)$ to $(0, 1)$ (Figure .)

Let $S \subseteq \mathbb{R}^3$ be the epigraph of X . At the origin o , the unique normal v to S is $(0, 1, 0)$, which points along the positive y -axis. In every section of S from v , the mean curvature of S at o is zero. The sectional curvature interval of S at o from v is then $[0]$. The smooth offset interval of S at o from NF^S is determined by C , and it is equal to $[-1, 1]$. As stated previously, a curved surface may have a sectional

curvature interval of $[0]$, but the smooth offset interval of a curved surface always contains a nonzero value.

CHAPTER 11

TIGHT EMBEDDINGS

11.1. Overview

We begin by briefly reviewing definitions of mean and Gaussian curvature obtained from the shape operator, which is the negative derivative of a surface's normal field. We define smoothly embedded surfaces, which are an important class of sets in our study of the smooth theory tight embeddings. We define a smoothly embedded surface as tight if and only if it minimizes total absolute Gaussian curvature out of all smoothly embedded surfaces isotopic to it. We then present a theorem by Chern and Lashof that imposes a lower bound on a surface's total absolute curvature proportional to the sum of its Betti numbers. We compare this to the Gauss-Bonnet theorem, which states that the integral of a surface's Gaussian curvature is equal to 2π times its Euler characteristic, which is the alternating sum of its Betti numbers.

We next consider tightness for three-dimensional polyhedra. We locate Gaussian curvature at polyhedron vertices, with the Gaussian curvature value at a vertex given by the angle defect formula (which we also discuss in Subsection 3.2.3.) We distinguish between the absolute value of the Gaussian curvature at a vertex, the sum of the positive and negative Gaussian curvature at a vertex, and the absolute Gaussian curvature at a vertex. We define a polyhedron as tight if and only if it minimizes the sum of its vertices' absolute Gaussian curvatures. We then present polyhedral versions of the Gauss-Bonnet and Chern-Lashof theorems.

We describe the equivalence in three dimensions of tightness, the two-piece property, and the minimization of height function maxima. While the two-piece property and minimization of height function maxima extend the definition of tightness from polyhedra and smoothly embedded surfaces to compact subsets of \mathbb{R}^3 , we show that

sets with the two-piece property do not necessarily minimize total absolute curvature in higher dimensions. We define the index of a height function critical point for a smoothly embedded surface as the number of negative principal curvatures at that point. If a smoothly embedded surface has a minimum quantity of critical points for each index, then it is tight, but we show that not all tight embeddings minimize the quantity of critical points for every index.

We close by adapting material from Kuiper [37] to define height function critical points for arbitrary sets, including critical points other than maxima. We define outward-pointing normals at critical points, and we explain the relationship between the quantity of height function critical points and slack.

11.2. Smooth and polyhedral tightness

11.2.1. Curvature from the shape operator.

Summary 17. We express the mean and Gaussian curvature of manifolds with differentiable normal fields in terms of the shape operator. The shape operator is the negative of the derivative of a normal field. At a point on a $(d - 1)$ -dimensional surface in \mathbb{R}^d , its value is a real symmetric matrix with $d - 1$ rows and columns. The Gaussian curvature is the determinant of the shape operator, and we take the mean curvature to be the trace of the shape operator divided by $d - 1$.

11.2.1.1. *The shape operator.*

Remark 11.1. In our definition of the shape operator, we omit details concerning the local intrinsic parametrization of a surface in the neighborhood of a surface point as well as the means of differentiating over that parametrization. Differential geometry monographs provide more rigorous treatments of these issues [35, 59].

Our presentation of curvature and tightness in this section focuses on what we call smoothly embedded surfaces:

Definition 11.1. A set $S \subseteq \mathbb{R}^d$ is a *smoothly embedded surface* if and only if S is a compact, oriented, and embedded $(d - 1)$ -dimensional manifold, and at every point p of S there is a unique unit normal that is differentiable over a local intrinsic parametrization of the neighborhood of p .

Notation 11. For a smoothly embedded surface $S \subseteq \mathbb{R}^d$ and $p \in S$, let $x \in \mathbb{R}^{d-1}$ intrinsically parametrize the neighborhood of point p . Let $N(x)$ denote the normal to S at the point corresponding to the parameter value x , and let $N(p)$ denote the normal to S at p .

Notation 12. For $x \in \mathbb{R}^d$ for some $d \in \mathbb{Z}^>$, let D_x denote differentiation with respect to x .

Definition 11.2. The *shape operator* for S , denoted σ^S , is the negative derivative of N with respect to x : $\sigma^S := -D_x N$.

Notation 13. The value of the shape operator at a point $p \in S$, denoted $\sigma^S(p)$, is the negative derivative of the normal field at p .

Remark 11.2. For a smoothly embedded surface $S \subseteq \mathbb{R}^d$, the value of the shape operator $\sigma^S(p)$ at a surface point $p \in S$ can be expressed by a real symmetric matrix M with $d - 1$ rows and columns. The matrix M represents how each of the $d - 1$ components of the normal vector to S at p vary as we move in $d - 1$ directions along S from p . Because M is real and symmetric, its eigenvectors are orthogonal, and its eigenvalues are real. The eigenvectors of M are the principal curvature directions for the surface S at p , while their corresponding eigenvalues are the principal curvatures at p .

11.2.1.2. *Gaussian curvature.*

Definition 11.3. For a smoothly embedded surface $S \subseteq \mathbb{R}^d$ and point $p \in S$, the *Gaussian curvature* of S at p , denoted $\gamma^S(p)$, is defined as

$$\gamma^S(p) := \det(\sigma^S(p))$$

Remark 11.3. The Gaussian curvature at p is the product of the principal curvatures at p , which is the determinant of M . As a determinant, the Gaussian curvature at p measures an “area change” that characterizes the spread of the normals around p . More precisely, it is the measure of the set of normals over a small neighborhood of p divided by the measure of the neighborhood. For instance, if the neighborhood is flat, it has only one normal, so its $(d - 1)$ -dimensional Hausdorff measure is zero. By contrast, the neighborhood of a point on a smoothly embedded surface in \mathbb{R}^d has a small, but positive, $(d - 1)$ -dimensional measure. The ratio of the measure of the normals for a flat neighborhood to the measure of the neighborhood is then zero, implying a zero Gaussian curvature for flat regions. For a unit sphere, the measure of the normals and the measure of the neighborhood are equal, corresponding to a Gaussian curvature of one. On a sphere in three dimensions of radius $\frac{1}{2}$, the measure of the normals is four times the measure of the neighborhood; more generally, a sphere of radius $\frac{1}{r}$ in three dimensions has a Gaussian curvature of r^2 .

11.2.1.3. Mean curvature.

Definition 11.4. For a smoothly embedded surface $S \subseteq \mathbb{R}^d$ and point $p \in S$, the *mean curvature* of S at p , denoted $\mu^S(p)$, is defined as

$$\mu^S(p) := \frac{\text{tr}(\sigma^S(p))}{d - 1}$$

Remark 11.4. The mean curvature at p is the mean of the principal curvatures at p . Some authors instead define the mean curvature at p as the trace of M , which is the sum of the principal curvatures at p . We can illustrate a difference between the two definitions using the unit sphere. If mean curvature is the mean of the principal curvatures, the mean curvature of the sphere is equal to one in any number of dimensions. If the mean curvature is instead the sum of the principal curvatures,

the sphere’s mean curvature is $d - 1$ in d dimensions. Higher-dimensional spheres “bend more” than lower-dimensional spheres in the sense that they bend in more directions, but we define mean curvature as the mean of the principal curvatures to preserve relationship between mean curvature and radius of curvature. Taking the mean of the principal curvatures ensures that if a d -dimensional sphere has mean curvature $\frac{1}{r}$, its radius of curvature is r in any finite dimension $d \in \mathbb{Z}^>$.

11.2.2. Tight smooth sets.

Summary 18. We define the total absolute curvature of a smoothly embedded surface as the integral over the surface of the absolute value of its Gaussian curvature. We define tight surfaces as minimizers of total absolute Gaussian curvature, then present a theorem by Chern and Lashof that places a lower bound on a surface’s total absolute curvature as a function of its topology. We contrast the Chern-Lashof Theorem with the Gauss-Bonnet Theorem, which states that a total Gaussian curvature smoothly embedded surface in three dimensions depends exclusively on its topology.

11.2.2.1. Definition of tight submanifolds.

Definition 11.5. The *total absolute curvature* of a smoothly embedded surface $S \subseteq \mathbb{R}^d$ is the integral $\int_S |\gamma^S|$ over S of the absolute value of its Gaussian curvature.

Definition 11.6. A smoothly embedded surface $S \subseteq \mathbb{R}^d$ is *tight* if and only if there is no smoothly embedded surface T isotopic to S such that $\int_T |\gamma^T| < \int_S |\gamma^S|$.

Remark 11.5. Tight surfaces minimize total absolute curvature for a given topology. Slack and total absolute curvature are identical for smoothly embedded surfaces, provided we compute slack using normals from only one side of the surface; see Theorem 11.7. Qualitatively, total absolute curvature measures the extent to which a surface folds over on itself when offset along its normals. The minimum amount of foldover for an embedding of a surface is determined by its topological type. Folds

are produced oscillations in the surface's normal field, so minimizing total absolute curvature minimizes those oscillations and their associated curvature critical points.

11.2.2.2. *Tightness and topological complexity.*

Notation 14. For a smoothly embedded surface S , let $b_i(S)$ denote the i^{th} Betti number of S .

Theorem 11.1. (*Chern and Lashof*) *If $S \subseteq \mathbb{R}^d$ is a smoothly embedded surface, then*

$$\int_S |\gamma^S| \geq \text{H}^{d-1}(\mathbb{S}^{d-1}) \left[\frac{1}{2} \sum_{i=0}^{\infty} b_i(S) \right]$$

Remark 11.6. This theorem asserts that a surface's topological complexity imposes a lower bound on its total absolute curvature. For example, a smooth embedding of the sphere \mathbb{S}^2 in \mathbb{R}^3 has a total absolute curvature of at least 4π , which is the surface area $\text{H}^2(\mathbb{S}^2)$ of the unit sphere multiplied by half the sum of the sphere's Betti number sequence $1, 1, 0, 0, \dots$. A smooth embedding of the torus in \mathbb{R}^3 has a total absolute curvature of at least 8π , which is twice as large as the bound for the sphere because the sum of the Betti number sequence $1, 2, 1, 0, 0, \dots$ for the torus is twice as large as the sum for the sphere.

Any convex smoothly embedded sphere in \mathbb{R}^3 has a total absolute curvature of exactly 4π , while a round torus provides an example of a torus embedding with a total absolute curvature of 8π . These embeddings necessarily minimize total absolute curvature, so they are tight under Definition 11.6.

11.2.2.3. *Tightness and the Gauss-Bonnet Theorem.*

Remark 11.7. We can compare Theorem 11.1 to the Gauss-Bonnet Theorem:

Notation 15. Let $\chi(T)$ denote the Euler characteristic of topological space T , given as the alternating sum of the Betti numbers of T :

$$\chi(T) = \sum_{i=0}^{\infty} (-1)^i b_i(T)$$

Theorem 11.2. (*Gauss-Bonnet*) *If $S \subseteq \mathbb{R}^3$ is a smoothly embedded surface, then*

$$\int_S \gamma = 2\pi\chi(S)$$

For any smooth embedding of sphere in three dimensions, $\int_S \gamma = 4\pi$. This value is equal to the total absolute curvature of a tight embedding of a sphere in three dimensions, but using Gaussian curvature rather than the absolute value of Gaussian curvature removes the geometric restrictions imposed by tightness on how the sphere may be embedded. For a smooth embedding of the torus in three dimensions $\int_S \gamma = 0$, while the total absolute curvature of a tightly embedded torus in three dimensions is 8π . Unlike a sphere, a torus has unequal total Gaussian curvature and total absolute curvature.

The mechanism behind the difference can be seen more simply in two dimensions, where we let κ denote the signed curvature of a smoothly embedded curve. Then for any smoothly embedded topological circle C , the two-dimensional analog of the Gauss-Bonnet Theorem states $\int_C \kappa = 2\pi$. Conventionally, we define signed curvature as

$$\kappa(s) = \frac{ds}{d\theta}$$

Where s is arclength and $\theta \in \mathbb{S}$ is direction of the normal at $C(s)$. If we pick point p on C and move along C in a counterclockwise direction, our normal makes exactly one counterclockwise rotation by the time we complete our traversal of C and return to p . The normal's counterclockwise rotation may reverse during the traversal, with clockwise normal rotation corresponding to a negative curvature, but the rotation direction necessarily reverses again and cancels the negative curvature in the integral. (Figure .)

The Gauss-Bonnet Theorem states that the integral of the signed normals to the sphere covers the sphere of normal directions exactly once. We use a surface of revolution to illustrate how a reversal of normal direction in two dimensions corresponds to a three-dimensional fold in a surface's normal offset (Figure .) In the example, a normal direction associated with the fold has a positive sign twice and a negative sign once, so it has a signed normal value of one in our integral. A round torus, by contrast, the negative Gaussian curvature at the hyperbolic points along the inner wall of the torus precisely cancels the positive Gaussian curvature from its outer wall, leading to an integral Gaussian curvature of zero.

Now we return to the plane and consider total absolute curvature. The integral $\int_C |\kappa|$ of the absolute value of curvature of C may be arbitrarily high, because the intervals over which the normal direction reverses and rotates clockwise rather than counterclockwise do not cancel intervals of counterclockwise rotation. The curves that minimize $\int_C |\kappa|$ are precisely those where the normal to C exclusively rotates counterclockwise as we traverse C counterclockwise. Those minimizers of total absolute curvature — the tight embeddings of the circle in the plane — are the boundaries of convex sets; any concavity would entail a reversal of normal direction.

Similarly, in three dimensions the value of $\int_S |\gamma|$ may be arbitrarily high, because surface patches with negative Gaussian curvature, where the patch flips as S folds over on itself when it is offset along its normals, do not cancel the positively signed, unflipped patches. Minimizers of total absolute curvature minimize the integral of negative Gaussian curvature. As the boundaries of convex sets have zero negative Gaussian curvature, they are the tight embeddings of the sphere.

11.2.3. Polyhedral Gaussian curvature.

Summary 19. We present a theory of a Gaussian curvature for polyhedra that enables us to formulate a polyhedral theory of tightness. Focusing on three dimensions, we locate a polyhedron's Gaussian curvature at its vertices, where its value is given by the angle defect formula. We distinguish between the absolute value of the Gaussian

curvature at a vertex, the sum of the positive and negative Gaussian curvature at a vertex, and the absolute Gaussian curvature at that vertex. These distinctions arise with polyhedra but not with smooth surfaces because the vertices of polyhedra are singular.

Remark 11.8. As described in Subsection 3.2.3 and Remark 10.24, the Gaussian curvature of a polyhedron $P \subseteq \mathbb{R}^3$ concentrates at its vertices, and its value at a vertex $v \in \partial P$ is given by

$$G_v^P = 2\pi - \sum_{i=1}^n \alpha_i$$

Where $\alpha_1, \alpha_2, \dots, \alpha_n$ is the sequence of angles between the circularly ordered edges incident on v . We can visualize this formula by considering a polytope $P \subseteq \mathbb{R}^3$ that is equal to the convex hull of a finite set of vertices $v \subseteq \partial P$. We assume P is three-dimensional in the sense that its Hausdorff dimension is equal to three or its affine hull is \mathbb{R}^3 . Each face of P has a unique normal that maps to a point on the unit sphere. When two faces are incident, an arc from a great circle connects the two points corresponding to the faces' normals. The union of all such arcs partitions the sphere into polygonal regions, with a one-to-one mapping between regions and vertices. Each point in a region mapped to vertex v corresponds to a normal to P at v . The smaller the solid angle at a vertex, the larger its associated polygonal region on the sphere (Figure .)

When P is polyhedral but not convex, each face still has a unique normal that maps to a point on the unit sphere, and we can draw arcs connecting those points. However, the arcs do not decompose the sphere into convex faces that can be traversed with a simple counterclockwise motion. Instead, concave polyhedron edges generate arcs that “go backwards,” introducing flipped regions on the sphere that correspond to negative Gaussian curvatures at vertices.

Gaussian curvature concentrates at polyhedron vertices in higher dimensions as well, and we can calculate it using a generalization of the angle defect formula. In \mathbb{R}^d , we compute the Gaussian curvature at a vertex v by taking the difference between the measure of a $(d - 2)$ -dimensional sphere and the angles between the adjacent $(d - 2)$ dimensional faces incident on v .

As we noted in Subsection 3.2.3, the absolute Gaussian curvature at a vertex v is not simply the absolute value of the Gaussian curvature at v [3]. Consider the following definitions, specialized for \mathbb{R}^3 :

Definition 11.7. A set $P \subseteq \mathbb{R}^3$ is a *polyhedron* if and only if it is a compact, topologically regular set such that there exists an $X \subseteq \wp(\wp(\mathbb{R}^3))$ with $\#X \in \mathbb{Z}^>$ such that for all $y \in X$ we have $\#y \in \mathbb{Z}^>$ and $\text{AH}(y) = \mathbb{R}^3$, with

$$P = \bigcup_{y \in X} \text{CH}(y)$$

Remark 11.9. We define a polyhedron as a topologically regular union of a finite number of convex hulls, where each convex hull is the convex hull of a finite number of points. We assume that the boundary of a polyhedron can be decomposed into faces, edges, and vertices.

Definition 11.8. For a polyhedron $P \subseteq \mathbb{R}^3$, the *Gaussian curvature* at a vertex v of P , denoted K_v^P , is given by the angle defect formula:

$$K_v^P := 2\pi - \sum_{i=1}^n \alpha_i$$

Where $\alpha_1, \alpha_2, \dots, \alpha_n$ is the sequence of angles between the circularly ordered edges incident on v .

Remark 11.10. We now provide definitions of positive, negative, and combined Gaussian curvature based on those given by Banchoff and Kühnel [3].

Definition 11.9. The *positive Gaussian curvature* at v , denoted K_v^{P+} is equal to zero if for all $r \in \mathbb{R}^>$, we have $v \in \text{CH}(P \cap B_r(v))$. Otherwise, K_v^{P+} is equal to the limit as $r \rightarrow 0$ of the Gaussian curvature at v of the convex hull of the intersection of the union E of the edges of P with a ball of radius r centered at v :

$$K_v^{P+} := \lim_{r \rightarrow 0} K_v^{\text{CH}(E \cap B_r(v))}$$

The *negative Gaussian curvature* at v , denoted K_v^{P-} , is defined as:

$$K_v^{P-} := K_v^{P+} - K_v^P$$

The *combined Gaussian curvature* at v , denoted K_v^{P*} , is defined as

$$K_v^{P*} := K_v^{P+} + K_v^{P-}$$

Remark 11.11. Both positive Gaussian curvature K_v^{P+} and negative Gaussian curvature K_v^{P-} are nonnegative at every vertex v . As a result, both $K_v^P \leq K_v^{P*}$ and $|K_v^P| \leq K_v^{P*}$: the Gaussian curvature at v and the absolute value of the Gaussian curvature at v are both less than or equal to the combined Gaussian curvature at v .

The combined Gaussian curvature at vertex v is not necessarily equal to the slack at v , nor does it match the mean Gaussian curvature of a slightly rounded version Q of P over a small neighborhood of the image of v on Q . For instance, let P consist of four thin triangular fins that meet at vertex v (Figure .) By making the angles from the fins that are incident on v very small, we can make the Gaussian curvature at v and the positive Gaussian curvature at v arbitrarily close to 2π , entailing a negative Gaussian curvature at v arbitrarily close to zero. The combined Gaussian curvature at v is then close to 2π .

When we compute the slack at v , we slightly dilate P in the neighborhood of v and count the number of connected components of points with each normal direction

over the neighborhood. Because P resembles a long, thin spike, we obtain a count of three for almost every direction on a hemisphere, leading to a slack very close to 6π .

Under our definition of a polyhedron, there exists a $\delta \in \mathbb{R}^>$ such that for all $\epsilon \in (0, \delta)$, the slightly rounded polyhedron $P \circ_\epsilon \bullet_\epsilon$ is ϵ -regular and isotopic to P . We can also choose ϵ to make the Hausdorff distance between P and $P \circ_\epsilon \bullet_\epsilon$ arbitrarily small. We adopt as a principle that the absolute Gaussian curvature at a vertex of P should match the mean absolute Gaussian curvature of $P \circ_\epsilon \bullet_\epsilon$ over a small geodesic neighborhood of the point $q \in \mathbb{R}^3$ corresponding to v on $P \circ_\epsilon \bullet_\epsilon$. We take q to be the closest point to v on $\partial(P \circ_\epsilon \bullet_\epsilon)$. This motivates a novel definition of polyhedral absolute Gaussian curvature:

Definition 11.10. For a polyhedron $P \subseteq \mathbb{R}^3$ with vertex v , let $S_\epsilon = \partial(P \circ_\epsilon \bullet_\epsilon)$. Then the *absolute Gaussian curvature* of P at v , denoted $K_v^{P||}$, is defined as

$$\lim_{r \rightarrow 0} \lim_{\epsilon \rightarrow 0} \frac{\int_{g_r^{S_\epsilon}(cp^{S_\epsilon}(v))} |\gamma^{S_\epsilon}|}{\mathbb{H}^{d-1}(g_r^{S_\epsilon}(cp^{S_\epsilon}(v)))}$$

Remark 11.12. Assume P is a long, thin spike composed of four fins. The two-dimensional absolute curvature of the boundary of the spike's intersection with a plane orthogonal to the spike axis is 6π (Figure .) Because the fin edges incident on v form small angles, we can integrate the two-dimensional absolute curvature over the hemisphere to yield an absolute Gaussian curvature close to 6π at v , which is equal to the slack of P at v .

In the spike cross-section, the integral of positive curvature is 4π , while the integral of the negative curvature is 2π . By contrast, the Gaussian curvature and positive Gaussian curvature are both 2π , implying zero negative Gaussian curvature. Both the positive and negative Gaussian curvatures are 2π less than the positive and negative curvatures in the planar cross section. More generally, taking the convex hull in computing positive Gaussian curvature reduces both positive and negative sectional curvature in the neighborhood of a vertex, which reduces both the positive and negative Gaussian curvature at the vertex. As a result, absolute Gaussian curvature is

greater than or equal to combined Gaussian curvature, which as we noted earlier is greater than or equal to the absolute value of Gaussian curvature: $|K_v^P| \leq K_v^{P*} \leq K_v^{P\parallel}$.

11.2.4. Tight polyhedral sets.

Summary 20. We define tight polyhedra as those that minimize absolute Gaussian curvature for their topological type. We then present a polyhedral form of the Gauss-Bonnet Theorem and compare it to Descartes' Lost Theorem. We present a theorem for polyhedra by Brehm and Kühnel that, like the theorem by Chern and Lashof, places a lower bound on total absolute curvature as a function of topological complexity. We close by noting that while the Gauss-Bonnet and Chern-Lashof theorems are essentially identical to their polyhedral counterparts, theorems concerning smooth and polyhedral tightness are not generally interchangeable.

Definition 11.11. The *total absolute Gaussian curvature* of a polyhedron $P \subseteq \mathbb{R}^3$ with vertex set V is the sum $\sum_{v \in V} K_v^{P\parallel}$ over the vertices in V of the absolute Gaussian curvature at each vertex.

Definition 11.12. A polyhedron $P \subseteq \mathbb{R}^3$ with vertex set V^P is *tight* if and only if there is no polyhedron $Q \subseteq \mathbb{R}^3$ with vertex set V^Q isotopic to P such that

$$\sum_{v \in V^Q} K_v^{Q\parallel} < \sum_{v \in V^P} K_v^{P\parallel}$$

Summary 21. Polyhedral versions of the Gauss-Bonnet Theorem and Theorem 11.1 by Chern and Lashof exist. We present them in three dimensions, starting with the polyhedral Gauss-Bonnet theorem:

Theorem 11.3. *For a convex polyhedron $P \subseteq \mathbb{R}^3$ with vertex set V ,*

$$\sum_{v \in V} K_v^P = 2\pi\chi(P)$$

Remark 11.13. We compare this to Descartes' Lost Theorem, which states that the sum of the interior angles over all faces of a convex polyhedron with v vertices is equal to $2\pi v - 4\pi$. If we let x denote the sum of the interior angles, the polyhedral form of the Gauss-Bonnet Theorem states $2\pi v - x = 4\pi$, so $x = 2\pi v - 4\pi$, which is Descartes' Lost Theorem. Although specialized to convex polyhedra, Descartes' theorem is remarkable because it anticipates the Gauss-Bonnet Theorem by over two centuries.

We now consider a three-dimensional polyhedral version of Theorem 11.1:

Theorem 11.4. (Brehm and Kühnel [7]) For a polyhedron $P \subseteq \mathbb{R}^3$ with vertex set V ,

$$\sum_{v \in V} K_v^{P*} \geq 2\pi (4 - \chi(P))$$

Remark 11.14. The differences in form between Theorems 11.1 and 11.4 are due to the fact that Theorem 11.4 is specialized to three dimensions. Because $K_v^{P*} < K_v^{P||}$, the Brehm-Kühnel theorem implies the total absolute curvature of a surface is bounded from below by its topological complexity. In three dimensions, each handle we add to a polyhedron decreases its Euler characteristic by 2, increasing the value of $2\pi (4 - \chi(P))$ by 4π . A polyhedral surface homeomorphic to a sphere has a polyhedral total absolute curvature of at least 4π , a surface with one handle has a total absolute curvature of at least 8π , and so on.

These lower bounds match the lower bounds for smoothly embedded surfaces, but theorems concerning smooth and polyhedral tightness are not trivially interchangeable. For instance, there exists a tight polyhedral immersion in \mathbb{R}^3 of the real projective plane with one handle, but there is no tight smooth immersion of the same manifold[8].

11.3. Minimizing height function maxima in \mathbb{R}^3 and \mathbb{R}^d

11.3.1. Zero tightness and the two-piece property.

Definition 11.13. For a choice of origin $o \in \mathbb{R}^d$, let the *height function* $f := \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be defined for vector $v \in \mathbb{R}^d$ and point $p \in \mathbb{R}^d$ as

$$f_v(p) := v \cdot (p - o)$$

Remark 11.15. Function f measures the distance of point p from the $(d - 1)$ -dimensional plane through the origin o orthogonal to v .

Definition 11.14. For a compact set $S \subseteq \mathbb{R}^d$, let $p \in S$ be a *maximal* point of S for v if and only if there is no point $q \in S$ such that $f_v(p) < f_v(q)$.

The *topset* of a compact set S for direction $v \in \mathbb{S}^{d-1}$, denoted T_v^S , is the set of all maximal points of S for v [37].

Theorem 11.5. (*Kuiper*) *The boundary of the convex hull a connected, compact set S is the union of the convex hulls of its topsets* [37]:

$$\partial\text{CH}(S) = \bigcup_{v \in \mathbb{S}^{d-1}} \text{CH} \left(\bigcup T_v^S \right)$$

Example 11.1. If $S \subseteq \mathbb{R}^3$ is the outer wall of a round torus, the topsets of S consist of two circles and an uncountable number of isolated points. Taking the convex hulls of the circles we obtain two disks, and the union of the disks with the remaining single-point topsets is the boundary of the convex hull of S (Figure .)

Remark 11.16. We introduce the concept of the local topset, which is a set of points where ∂S has a local height function maximum:

Definition 11.15. The *local topset* of a compact, connected set $S \subseteq \mathbb{R}^d$ for direction $v \in \mathbb{S}^{d-1}$, denoted LT_v^S , is any set of points $t \in \{\{p \in S \mid f_v(p) = c\} \mid c \in \mathbb{R}\}$ of points at a fixed height for direction v such that there exists a $\delta \in \mathbb{R}^>$ such that for all $\epsilon \in (0, \delta)$ there is no point $q \in g_\epsilon^S(t)$ such that for $p \in t$ we have $f_v(p) < f_v(q)$.

Remark 11.17. Departing from prior art, we define 0-tightness in terms of local topsets:

Definition 11.16. A compact set $S \subseteq \mathbb{R}^d$ is *0-tight* if and only if there is no set $T \subseteq \mathbb{R}^d$ isotopic to S such that

$$\sum_{i=1}^{\infty} \mathbb{H}^{d-1} \left(\left\{ v \in \mathbb{S}^{d-1} \mid i \leq \#\kappa \left(\bigcup \text{LT}_v^T \right) \right\} \right) < \sum_{i=1}^{\infty} \mathbb{H}^{d-1} \left(\left\{ v \in \mathbb{S}^{d-1} \mid i \leq \kappa \left(\bigcup \text{LT}_v^S \right) \right\} \right)$$

Remark 11.18. A set is 0-tight if and only if it minimizes the total quantity of height function maxima over all height functions. The term 0-tight, rather than tight, refers to the fact that 0-tight smoothly embedded surfaces have a minimal quantity of critical points of index zero. The definition's use of height function maxima rather than curvature broadens its scope by making it applicable to sets with nondifferentiable normals, where the shape operator and the curvatures obtained from it are undefined. For instance, a closed line segment in \mathbb{R}^3 is 0-tight under this definition, but its shape operator is undefined. We can envision oriented planes rolling over the segment, with each plane normal to v that contains the segment yielding a value $\#\kappa(\text{LT}_v^S) = 1$. The planes through the segment endpoints whose normals point away from the segment interior also yield $\#\kappa(\text{LT}_v^S) = 1$. Under the Gauss map, each segment endpoint maps to an open hemisphere, while the segment itself maps to a great circle (Figure .)

The property of 0-tightness is closely related to the two-piece property [3, 37]:

Definition 11.17. A compact set $S \subseteq \mathbb{R}^d$ has the *two-piece property* if and only if for any $(d - 1)$ -dimensional affine subspace H of \mathbb{R}^d , the set S/H consists of no more than two components: $\#\kappa(S/H) \leq 2$.

Theorem 11.6. *If a compact set $S \subseteq \mathbb{R}^d$ has the two-piece property then S is 0-tight [3, 37].*

PROOF. Suppose that there exists a direction $v \in \mathbb{S}^{d-1}$ such that the height function f_v has more than one maximum on S ; equivalently, the local topset LT_v^S is not connected. We can order the connected components $\kappa(\text{LT}_v^S)$ by their height function values x_1, x_2, \dots with x_1 being the greatest. Let $A \subseteq \mathbb{R}^d$ be a component with

height value x_1 , and let B be a component with height value x_2 . Consider the $(d - 1)$ -dimensional affine subspace H orthogonal to v positioned so that $B \subset H$. Because the points in B are local height function maxima, we can select an $\epsilon \in \mathbb{R}^>$ such that A and B lie in different connected components of $S / (H - \epsilon v)$. We conclude that if f_v has more than one maximum on S , then S does not have the two-piece property.

A nonempty set must have at least one height function maximum for every direction, and by the argument above a set with the two-piece property cannot have more than one. If S has the two-piece property, it has exactly one height function maximum for every direction, implying S is 0-tight. \square

Corollary 11.1. *If a compact set S has the two-piece property, then*

$$\sum_{i=1}^{\infty} \mathbb{H}^{d-1} \left(\left\{ v \in \mathbb{S}^{d-1} \mid i \leq \kappa \left(\bigcup \text{LT}_v^S \right) \right\} \right) = \mathbb{H}^{d-1} (\mathbb{S}^{d-1})$$

PROOF. Following the argument above, $\kappa \left(\bigcup \text{LT}_v^S \right) = 1$ for all $v \in \mathbb{S}^{d-1}$. The summation on the right hand side of the equation then only has one term: $\mathbb{H}^{d-1} (\{v \in \mathbb{S}^{d-1} \mid 1 \leq 1\})$. The corollary follows trivially from $\mathbb{S}^{d-1} = \{v \in \mathbb{S}^{d-1} \mid 1 \leq 1\}$. \square

Corollary 11.2. *If a compact set S has the two-piece property, then the local topsets of S lie on the boundary of the convex hull of S : for all $v \in \mathbb{S}^{d-1}$, $\text{LT}_v^S \subseteq \partial \text{CH}(S)$.*

PROOF. By Theorem 11.5, the boundary of the convex hull of a set is the union of the convex hulls of the topsets of S . This implies that the topsets of S lie on the boundary of the convex hull of S . Because a compact set S with the two-piece property has one height function maximum for every direction, every local topset of S is a topset S . Consequently, the local topsets of S lie on the boundary of the convex hull of S . \square

Remark 11.19. Under our definition of 0-tightness, there exist 0-tight sets that do not have the two-piece property. For instance, there exists a 0-tight embedding of the trefoil knot in \mathbb{R}^3 that does not have the two-piece property (Figure;) we omit the

details. In this respect we differ from Kuiper [37] and Banchoff and Kühnel [3], for whom 0-tightness and the two-piece property are equivalent. They require that the number of height maxima for a 0-tight set equal the zero Betti number b_0 , which is equal to one for a connected set.

In three dimensions, a smoothly embedded 0-tight surface $S \subseteq \mathbb{R}^3$ minimizes total absolute curvature, making it tight under Definition 11.6. The height function maxima are precisely those points with two positive principal curvatures, so they have positive Gaussian curvature. Given that the difference between the positive Gaussian curvature and negative Gaussian curvature is a constant by the Gauss-Bonnet Theorem, minimizing the amount of positive Gaussian curvature necessarily minimizes the amount of negative Gaussian curvature. It also minimizes absolute Gaussian curvature, which is the sum of both quantities. A similar argument shows that a 0-tight polyhedron $P \subseteq \mathbb{R}^3$ is tight.

11.3.2. Zero-tightness in \mathbb{R}^d and k -tightness.

Remark 11.20. A tangent plane to a smoothly embedded surface identifies a height function critical point, which is not necessarily a local height function maximum. Because the surface is smooth, its principal curvatures are defined, and we refer to the number of negative curvatures as the index of the height function critical point:

Definition 11.18. For a smoothly embedded surface $S \subseteq \mathbb{R}^d$, the *index* $p \in \partial S$, denoted $i(p)$ is the number of negative principal curvatures at p .

Remark 11.21. In higher dimensions 0-tightness does not imply tightness. For a smoothly embedded surface $S \subseteq \mathbb{R}^5$, we may have critical points with indices of 0, 1, 2, or 3. The Gaussian curvature of S is positive at a point with an even index, and negative at a point with an odd index. Suppose S is 0-tight. This requires that S minimize the measure of height function maxima, which have index 0. The generalized Gauss-Bonnet Theorem [cite] requires that the amount of positive Gaussian curvature over S minus the amount of negative Gaussian curvature over S equal a constant, but

we may make the quantity of height function maxima of index 1 or 3 arbitrarily large and then inflate the quantity of critical points of index 2 and still satisfy the theorem. This yields an S that is 0-tight, but has arbitrarily large total absolute curvature.

We can extend the notion of minimizing critical points of index 0 to that of minimizing critical points of index less than or equal to k :

Definition 11.19. A smoothly embedded surface $S \subseteq \mathbb{R}^d$ is k -tight if and only if there is no $T \subseteq \mathbb{R}^d$ isotopic to S such that for some $j \leq k$

$$\begin{aligned} & \sum_{i=1}^{\infty} H^{d-1}(\{v \in \mathbb{S}^{d-1} \mid i \leq \#\kappa(\{p \in \partial T \mid (N(p) = v) \wedge (i(p) = j)\})\}) \\ & \leq \sum_{i=1}^{\infty} H^{d-1}(\{v \in \mathbb{S}^{d-1} \mid i \leq \#\kappa(\{p \in \partial S \mid (N(p) = v) \wedge (i(p) = j)\})\}) \end{aligned}$$

Remark 11.22. Our definition of k -tightness differs from the definitions given by Kuiper [37] and Banchoff and Kühnel [3] in the same way that our definition of 0-tightness differs from theirs. We require that a k -tight set S minimize critical points of a given index k or less over all sets isotopic to S , while they require that the number of height function critical points with index i equal the Betti number b_i . Our definition is less restrictive.

If a smoothly embedded surface $S \subseteq \mathbb{R}^d$ is k -tight for all k in $0, 1, \dots, d-1$, then it is tight: each of its critical points is necessary, so its total quantity of critical points is minimal. This does not imply that an embedding that minimizes total absolute curvature, making it tight under Definition 11.6, is also $(d-1)$ -tight. We ask:

QUESTION: If $S \subseteq \mathbb{R}^d$ is a tight smoothly embedded surface, is S $(d-1)$ -tight?

Claim 11.1. No, a tightly embedded smooth surface $S \subseteq \mathbb{R}^d$ need not be $(d-1)$ -tight. We limit ourselves to outlining an argument for the claim. The “round” embedding $S \subseteq \mathbb{R}^5$ of the 4-torus $\mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{S}^1$ in \mathbb{R}^5 rotates $(\mathbb{S}^1)^k$ for $k = 1, 2, 3$ successively about three orthogonal axes $\alpha_1, \alpha_2, \alpha_3$ spaced to prevent self-intersection.

The result has total absolute curvature $\int_S |\gamma| = \frac{64\pi^2}{3}$. The sum of the Betti numbers for the 4-torus is 16, while the 4-dimensional Hausdorff measure of the 4-dimensional unit sphere $\mathbb{S}^4 \subseteq \mathbb{R}^5$ is $\frac{8\pi^2}{3}$. According to Theorem 11.1, the minimum possible total absolute curvature for the 4-torus is half the product of these two numbers, or $\frac{64\pi^2}{3}$, which implies that the round embedding minimizes total absolute curvature. However, this 4-torus is not 4-tight, because with a small deformation we may decrease the quantity of critical points of index 1 while increasing the quantity of critical points of index 3 by the same amount, preserving the quantity of points with index 0 and index 2. This keeps both the total negative and total absolute Gaussian curvature constant, but demonstrates that our original embedding was not 1-tight because the new embedding has a smaller quantity of critical points of index 1. Although S is tight in the sense that it minimizes total absolute curvature, it is not $(d - 1)$ -tight.

In the next section, we formally define height function critical points and their relationship to total absolute curvature.

11.4. Height function critical points

11.4.1. Tightness from height function critical points.

Summary 22. Given that the two-piece property and the minimization of height function maxima allow us to define 0-tightness for a large class of subsets of Euclidean space, we ask whether we can define tightness for arbitrary subsets of Euclidean space by using all height function critical points, rather than only maxima. We introduce a definition of a height function critical point for an arbitrary set, then modify the definition to obtain outward-pointing normals to a set. We illustrate a difference between the normals we define using height function critical points and the normals we define in Section 10.3 using morphological operations. We then compare slack to the measure of the set of height function critical points. We show that the two measures are equivalent on smoothly embedded surfaces, but slack is equivalent to

total absolute curvature on polyhedra while the measure of height function critical points is not.

11.4.2. Critical points and normals.

Remark 11.23. To test for the presence of a height function critical point along direction $v \in \mathbb{S}^{d-1}$ at point $p \in \mathbb{R}^d$, we examine a small cylinder centered at p with its axis aligned with v . The following notation allows us to concisely specify a cylinder:

Notation 16. For $p, q \in \mathbb{R}^d$ and $r \in \mathbb{R}^>$, let $C_r(\overline{pq})$ denote the trimmed right circular cylinder with axis \overline{pq} and radius r :

$$C_r(\overline{pq}) := \{x \in \mathbb{R}^d \mid (cp^{\overline{pq}}(x) \in \overline{pq}/p/q) \wedge (d(x, \overline{pq}) < r)\}^-$$

Remark 11.24. We define $p \in \mathbb{R}^d$ as a height function critical point of $S \subseteq \mathbb{R}^d$ along $v \in \mathbb{S}^{d-1}$ if and only if there is no isotopy over the interval of cross-sections of S orthogonal to v in a small cylinder centered on p :

Definition 11.20. For $S \subseteq \mathbb{R}^d$, point $p \in \mathbb{R}^d$, and vector $v \in \mathbb{S}^{d-1}$, let ϖ be the $(d-1)$ -dimensional affine subspace containing p and orthogonal to v . Then p is a *critical point* for S along direction v if and only if there exists an $r \subseteq \mathbb{R}^>$ such that for all $\rho \in (0, r)$, there exists a $\delta \in \mathbb{R}^>$ such that for all $\epsilon \in [0, \delta]$, there is no isotopy $I : [-\epsilon, \epsilon] \times \mathbb{R}^{d-1} \rightarrow \mathbb{R}^{d-1}$ where for all $t \in [-\epsilon, \epsilon]$,

$$I_t = S \cap C_\rho \left(\overline{(p - \epsilon v)(p + \epsilon v)} \right) \cap (\varpi + tv)$$

Notation 17. For $S \subseteq \mathbb{R}^d$ and $v \in \mathbb{S}^{d-1}$, let CP_v^S denote the set of critical points for S along direction v .

Remark 11.25. We define direction $v \in \mathbb{S}^{d-1}$ as a normal to $S \subseteq \mathbb{R}^d$ at $p \in \mathbb{R}^d$ if and only if p is a critical point for S along v and v does not “point into” S , in the

sense that there is no open interval with an endpoint at p extending in the direction of v that is contained in S :

Definition 11.21. For $S \subseteq \mathbb{R}^d$, point $p \in \mathbb{R}^d$, and vector $v \in \mathbb{S}^{d-1}$, vector v is a *normal* to S at p if and only if for all $r \in \mathbb{R}^>$, the open line segment $\overline{p(p+rv)}/p/(p+rv) \not\subseteq S$.

Remark 11.26. We now illustrate a difference between the normals defined in Section 10.3 and the normals from Definition 11.21. The normals to a set $S \subseteq \mathbb{R}^d$ at $p \in \partial S$ as defined in Section 10.3 are independent of which subset of the boundary of a connected component k of S or S^c in the neighborhood of p is contained in k . This is because in Definition 10.9 we take the closure of k before testing whether a candidate normal to S at p is normal to k^- . There is no comparable closure mechanism in Definition 11.21, so adding or removing points to the boundary of S can affect its normal field, even when $(\partial S)^\circ = \emptyset$. Consider the following example:

Example 11.2. Let $H \subseteq \mathbb{R}^2$ be a half-plane. The height function critical point normals from Definition 11.21 and the normals defined in Section 10.3 both place a normal at each point on ∂H orthogonal to ∂H and pointing into the exterior of H . For $p \in \partial H$, the height function critical point normals to $H \cup \{p\}$ include a fan of normals subtending an arc of $\frac{\pi}{2}$, pointing toward the semicircle $\partial B_1(p) \cap H^c$ (Figure .) By contrast, the normals from Section 10.3 to $H \cup \{p\}$ are identical to the normals to H : $\text{NF}^H = \text{NF}^{H \cup \{p\}}$.

11.4.3. Tightness from critical points and comparison to slack.

Remark 11.27. To measure the total quantity of height function critical points for a set $S \subseteq \mathbb{R}^d$, we integrate over the sphere the number of connected components of critical points for each direction:

Definition 11.22. For $S \subseteq \mathbb{R}^d$, the *total critical point measure* of S , denoted TCM^S , is defined as

$$\text{TCM}^S := \sum_{i=1}^{\infty} \mathbb{H}^{d-1} (\{v \in \mathbb{S}^{d-1} \mid \#\kappa(\text{CP}_v^S) \leq i\})$$

Remark 11.28. Our definition of total critical point measure closely resembles our definition of slack. It omits the means for measuring subsets of critical points as well as any reference to normal count. For smoothly embedded surfaces, the normal count is always one, and we claim that total absolute curvature, total critical point measure, and slack are equivalent. More precisely, the total critical point measure and slack of a smoothly embedded surface are equal to twice its total absolute curvature:

Theorem 11.7. *For a smoothly embedded surface $S \subseteq \mathbb{R}^d$,*

$$\text{TCM}^S = \text{SL}^S = 2 \int_S |\gamma^S|$$

PROOF. The normal count at every point $p \in \partial S$ is equal to one. The normal count for direction v as given by Definition 10.13 then reduces to the number of connected components of S that have v as a normal. Because S is a closed set with no interior, point p is a critical point for direction v under Definition 11.20 if and only if v is a normal to S at p under Definition 10.9. Slack, which is the integral over all directions of the sum of normal counts for each direction, is then equal to the total critical point measure, which is the integral over all directions of the number of connected components of critical points for S for each direction.

Slack and total absolute curvature both measure the “size” of a set’s normal field, or the total variation of the normals of a set over its boundary. Slack integrates the number of connected components of points with a given direction over all possible directions. Total absolute curvature integrates the measure of the set of normals in the neighborhood of a surface point over all the points on a surface. They compute the same quantity for smoothly embedded surfaces up to a scaling factor of two because slack’s integration of points over directions mirrors total absolute curvature’s integration of normal variation over surface points. The factor of two difference

between them is because slack by default uses normals from both sides of an oriented surface while total absolute curvature uses normals from only one side. \square

Remark 11.29. In Subsection 11.2.3, we described a polyhedron $P \subseteq \mathbb{R}^3$ such that $\text{TCM}^P < \sum_{v \in V} K_v^{P||}$. We now show $\text{SL}_v^P = K_v^{P||}$ at vertex v of any polyhedron P .

Theorem 11.8. *For polyhedron $P \subseteq \mathbb{R}^d$ with vertex v , the slack at v is equal to the absolute Gaussian curvature at v : $\text{SL}_v^P = K_v^{P||}$.*

PROOF. [Work in progress — this is illustrative of the idea, but incomplete.]

We compute slack at v from the normal counts over the neighborhood of v . The normal count for a given normal direction x is the number of connected components of points from a slight dilation or erosion of the neighborhood of v that include x as a normal. We compute absolute Gaussian curvature using small-radius morphological opening and closing, rather than dilation or erosion. We argue that as the radius $\epsilon \in \mathbb{R}^>$ we use goes to zero, the number of connected components from the dilation or erosion of the neighborhood of v that have a given normal matches the number of connected components with that normal from the result of opening and closing the neighborhood of v .

Each $(d - 2)$ -dimensional edge of P incident on v has a $(d - 2)$ -dimensional fan of normals. Over a neighborhood of v , a subset of the edges will have a given normal $x \subseteq \mathbb{S}^{d-1}$. However, those edges are connected at v . We quantify the number of times a normal is “used” by blending the neighborhood of v to break the set of points with normal x into separate connected components, which we then count.

If v is locally convex, then dilation will replace it with a spherical cap. Dilation replaces the convex edges with cylindrical patches, keeping concave edges sharp. For a sufficiently small dilation radius, the fan of normals across a cylindrical patch remains identical to the fan of normals at a point on a convex edge, while the fan of normals at a point on a concave edge is identical before and after dilation. The connected

components of the dilated neighborhood of v with normal x are still lines, but the lines stop at the spherical cap. The number of lines is the normal count.

If instead of dilating the neighborhood of v we open and close it, the result is smooth, with both concave and convex edges blended. For a small blending radius, the fan of normals across a blended edge is identical to the fan across the edge itself. As with the dilated neighborhood of v , the set of points with normal x is a set of lines that stop near a spherical cap at v . Whether we dilate the neighborhood of v or open and close it, the number of lines with normal direction x is equal to the number of edges incident on v with normal direction x . A similar analysis shows that when v is locally concave, the erosion of the neighborhood of v has the same number of connected components of points with normal x as the result of opening and closing the neighborhood of v .

If v is a saddle vertex — contained in the convex hull of its incident edges — then dilation maps the convex edges to cylindrical patches that meet at $(d - 2)$ -dimensional arcs that smoothly join to the $(d - 2)$ -dimensional offsets of the concave edges incident on v . The connected components of points with normal direction x stop at the arcs. As at a locally convex or concave vertex, the number of connected components with normal x is equal to the number of edges with normal x .

If we open and close the neighborhood of saddle vertex v , we obtain a complex blend at v , joining convex and concave cylindrical patches that correspond to the convex and concave edges incident on v . The connected components of points with normal x are line segments along those cylindrical patches, and as before the number of segments is equal to the number of edge incident on v with normal x . The blending patch near v keeps the segments from joining into a single connected component.

Because the number of connected components of points with normal direction x is equal in our computations of both the slack and absolute Gaussian curvature at v , we conclude $\text{SL}_v^P = \text{K}_v^{P||}$. □

Remark 11.30. The theory of tight embeddings has historically addressed problems such as constructing a tight embedding of a manifold or determining which manifolds can be tightly embedded or immersed. In the next chapter, we adapt ideas from the theory of tight embeddings to problems involving containment constraints by introducing a variant of the two-piece property called the hull property. We use the hull property to define abstract hulls, which captures the idea of wrapping a surface around a set or separating two sets with a membrane-in-tension.

CHAPTER 12

COVERS AND HULLS

12.1. Hulls and relative hulls

Summary 23. We introduce hulls, which capture the abstract notion of “wrapping” a set with an enclosing set. The Euclidean convex hull, the relative convex hull, and the geodesic hull are all hulls under our definition of a hull. However, our definition relates more strongly to the theory of tight embeddings than to the theory of convex structures. Echoing the relationship between 0-tightness and the two-piece property, we formulate our hull definition in terms of how a plane cuts a set into pieces. Rather than requiring a set to yield at most two connected components for each possible cut, which is the necessary and sufficient condition for 0-tightness, we require that when a plane cuts a hull $S \subseteq \mathbb{R}^d$ of $R \subseteq S$, each of the resulting connected components intersects R . When R is connected, this entails that a plane cannot “shave off” a piece of S without intersecting R .

We first prove that Euclidean convex hulls are hulls (Theorem 12.1.) We then prove the Shaving Lemma (Theorem 12.2,) which asserts that a set $S \subseteq \mathbb{R}^d$ is a hull of $R \subseteq S$ if and only if every height function maximum on the boundary of S is also a height function maximum on R . If S is a hull, the Shaving Lemma implies it has no “bumps” or locally convex points that are not also locally convex points of R , because a cutting plane can shave off any such bumps. Conversely, if it is impossible to shave any bumps off S , then S is a hull. We next introduce minimal hulls, where a minimal hull of R minimizes boundary measure over all hulls of R , under conditions we specify. The geometric properties of minimal hulls differ markedly from those of convex hulls, indicating the range of sets classified as hulls under our hull definition.

Next, we define relative hulls, which contain one set $R \subseteq \mathbb{R}^d$ and exclude a second set $G \subseteq \mathbb{R}^d$. The difference between hulls and relative hulls is that while a connected component of the complement of a hull cut by a plane is necessarily unbounded, a connected component of the complement of a relative hull cut by a plane may be unbounded or may contain a component of G . We prove results for relative hulls analogous to those we prove for hulls: we prove that relative convex hulls are hulls, then prove the Relative Shaving Lemma, and finally prove that relative minimal hulls are hulls. The proofs for relative hulls are technically and conceptually similar to those for hulls.

We discuss relative minimal hulls in relation to our earlier work on tightening, where we proposed that they serve as a generalization of the two-dimensional relative convex hull to higher dimensions. Although relative minimal hulls are hulls, they are not maximally convex. This motivates our definition of tight hulls in Chapter 13.

Definition 12.1. A set $H \subseteq \mathbb{R}^d$ is a $(d - 1)$ -plane if and only if H is a $(d - 1)$ -dimensional affine subspace of \mathbb{R}^d .

Remark 12.1. The following definition specifies an abstract meaning for the term “hull.” Informally, it states that for $S \subseteq \mathbb{R}^d$ and $R \subseteq S$, the set S is a hull of R if the boundary of S “wraps around” R with no “extraneous bumps” disjoint from R can be cut off by a $(d - 1)$ -plane. More precisely, the “extraneous bumps” are the height function maxima addressed by the Shaving Lemma (Theorem 12.2.) The requirement that $(\partial S)^\circ = \emptyset$ does not require that $(\partial R)^\circ$; for instance, the unit disk is a hull of the set of points with rational coordinates lying in the unit disk. Requiring $(\partial S)^\circ = \emptyset$ implies that the complete graph on the set of points with rational coordinates in a unit cube $[0, 1] \times [0, 1] \times [0, 1]$ is not a hull of itself. Instead, its hull consists of the unit cube and an arbitrary subset of the cube’s boundary.

possible to eliminate empty boundary interior requirement?

Definition 12.2. For a set $R \subseteq \mathbb{R}^d$, set $S \subseteq \mathbb{R}^d$ such that $R \subseteq S$ and $(\partial S)^\circ = \emptyset$ is a *hull* of R if and only if for every $(d-1)$ -plane $H \subseteq \mathbb{R}^d$, every connected component of S/H intersects R and every connected component of S^c/H is unbounded.

A set has the *hull property* if and only if it is a hull.

Remark 12.2. Although it is widely known that the boundary of a Euclidean convex hull has an empty interior, we include a proof of that property for completeness. The proof hinges on the fact that the convex hull of a set $S \subseteq \mathbb{R}^d$ is equal to the union of the convex hulls of all the subsets of S .

Lemma 12.1. For all $S \subseteq \mathbb{R}^d$, the interior of the boundary of the convex hull of S is empty: $(\partial \text{CH}(S))^\circ = \emptyset$.

PROOF. If $(\partial \text{CH}(S))^\circ \neq \emptyset$, let $b \in \mathbb{B}^{(\partial \text{CH}(S))^\circ}$ be a ball in the interior of the boundary of the convex hull of S such that $b \not\subseteq \text{CH}(S)$. Because the convex hull of S is convex, it contains the convex hull of each of the subsets of S :

$$\text{CH}(S) = \bigcup_{X \subseteq S} \text{CH}(X)$$

In particular, $\text{CH}(b \cap S) \subseteq \text{CH}(S)$. Because $b \subset \partial \text{CH}(S)$, we have $b \cap S \neq \emptyset$. If $b \not\subseteq \text{CH}(b \cap S)$, there then exists $p \in \text{CH}(b \cap S)$ and $q \in \text{CH}(S) \setminus \text{CH}(b \cap S)$ such that the closed line segment \overline{pq} is not contained in $\text{CH}(S)$. This implies the convex hull of S is not convex, which is a contradiction. If instead $b = \text{CH}(b \cap S)$, then $b \subseteq \text{CH}(S)$; in particular, $b \subseteq \text{CH}(S)^\circ$ because b is an open ball. Because $b \subseteq \partial \text{CH}(S)$ and $\text{CH}(S)^\circ \cap \partial \text{CH}(S) = \emptyset$, this is also a contradiction. We conclude that the interior of the boundary of a convex hull is empty. \square

Remark 12.3. We now prove that every convex hull is a hull. The essence of the proof is that if we can shave a piece from the convex hull, it necessarily remains convex but becomes a smaller set, contradicting that the convex hull is the smallest convex set containing S . If instead we can shave a bounded component from the

convex hull's complement — thereby identifying a dent in the hull — the convex hull cannot be convex because it does not contain every line segment that connects a pair of its points. It is then impossible for a convex hull to not be a hull, so every convex hull is a hull.

Theorem 12.1. *For any set $S \subseteq \mathbb{R}^d$, the convex hull of S is a hull of S .*

PROOF. For all $S \subseteq \mathbb{R}^d$, $S \subseteq \text{CH}(S)$ because the Euclidean convex hull defines a convex structure, and in every convex structure the convex hull of a set contains the set. By Lemma 12.1, $(\partial\text{CH}(S))^\circ = \emptyset$.

$$\overline{pq}$$

If there exists a $(d - 1)$ -plane H such that a connected component $k \in \kappa(\text{CH}(S) \setminus H)$ does not intersect S , then $\text{CH}(S) \setminus k$ must be convex. Otherwise, there exist points $p, q \in \text{CH}(S) \setminus k$ such that $\overline{pq} \not\subseteq \text{CH}(S) \setminus k$, but because p and q lie in $\text{CH}(S)$, segment $\overline{pq} \subseteq \text{CH}(S)$. Segment \overline{pq} then intersects k , but neither p nor q lies in k . Consequently, $\overline{pq} \cap (\partial k \setminus H) \neq \emptyset$, because a line segment not lying in H cannot intersect the planar patch $\partial k \cap H$ twice. Because $\partial k \setminus H \subseteq \partial\text{CH}(S)$, we can conclude that $\overline{pq} \not\subseteq \text{CH}(S)$, which contradicts that $\text{CH}(S)$ is convex. Therefore $\text{CH}(S) \setminus k$ is convex. However, $\text{CH}(S) \setminus k \subset \text{CH}(S)$, which contradicts the fact that $\text{CH}(S)$ is the smallest convex set containing S when all such sets are ordered by set inclusion. We conclude that every connected component of $\text{CH}(S) \setminus H$ intersects S .

If there is a $(d - 1)$ -plane H such that a connected component $k \in \kappa(\text{CH}(S)^c \setminus H)$ is bounded, then there exist points $p, q \in \text{CH}(S)$ such that the closed line segment \overline{pq} intersects k . Because $k \subseteq \text{CH}(S)^c$, this implies $\overline{pq} \not\subseteq \text{CH}(S)$. Then the convex hull of S is not convex, which is a contradiction. We conclude that every connected component of $\text{CH}(S)^c \setminus H$ is unbounded. The convex hull of S is then a hull of S . \square

Remark 12.4. The Shaving Lemma resembles the two-piece property. While the two-piece property requires a set to have one height function maximum for every

direction, the Shaving Lemma requires that each strict height function maximum over the boundary of a hull S of R coincide with a strict height function maximum of R^- . Informally, it requires that a hull of R be “stretched over” or “supported by” extreme points from R .

We begin by defining strict local height function maxima. Our proof then demonstrates that if there is a strict height maximum on the boundary of $S \subseteq \mathbb{R}^d$ of a hull of $R \subseteq S$ that is not a strict height maximum on R^- , then there exists a $(d - 1)$ -plane that cuts off a component of S or S^c disjoint from R , so S is not in fact a hull. Then next part of the proof shows that if S is not a hull of R and so we can shave off a piece of S or S^c , the boundary of the piece has a strict height function maximum that lies on ∂S but is not a strict height maximum on R^- . Consequently, if S is a hull of R , the strict height function maxima over its boundary are strict height function maxima of R^- . Conversely, if S has no strict height maxima over its boundary that are not strict height function maxima of R^- , then S is a hull of R .

Definition 12.3. For set $S \subseteq \mathbb{R}^d$, point $p \in S$, and direction $v \in \mathbb{S}^{d-1}$, point p is a *strict local maximum* of the height function f_v on S if and only if there exists a $\delta \in \mathbb{R}^>$ such that for all $\epsilon \in (0, \delta)$, for every point $q \in g_\epsilon^S(p)$ we have $f_v(q) < f_v(p)$.

Theorem 12.2. (*The Shaving Lemma*) For $S \subseteq \mathbb{R}^d$ such that $R \subseteq S$ and $(\partial S)^\circ = \emptyset$, the set S is a hull of R if and only if for every point $p \in \partial S$ and direction $v \in \mathbb{S}^{d-1}$ such that p is a strict local maximum of height function f_v over S^- , point p is a strict local maximum of f_v over R^- , and there is no $p \in \partial S$ and $v \in \mathbb{S}^{d-1}$ such that p is a strict local maximum of f_v over S^{c-} .

PROOF. Suppose p is a strict local maximum on S^- , but not a strict local maximum on R^- . The value of $f_v(p) = c$ for some $c \in \mathbb{R}$. Because p is a strict local maximum, there exists an $\epsilon \in \mathbb{R}^>$ such that for all $q \in g_\epsilon^{\partial S}(p)$, we have $f_v(q) < f_v(p)$. Therefore there exists an $a \in \mathbb{R}$ such that

$$a = \bigwedge \{x \in \mathbb{R} \mid \forall q \in g_\epsilon^{\partial S}(p) : f_v(q) > x\}$$

Because p is not a strict local maximum of f_v on R^- , there exists a $b \in \mathbb{R}$ such that $a \leq b \leq c$ and for the $(d-1)$ -plane $H = \{q \in \mathbb{R}^d \mid f_v(q) = b\}$, there is a bounded connected component $k \in \kappa(S \setminus H)$ such that $k \cap R = \emptyset$. Then S is not a hull because there exists a connected component of $S \setminus H$ disjoint from R .

Instead suppose p is a strict local maximum of f_v on S^{c-} . As before, there exists an $a \in \mathbb{R}$ such that

$$a = \bigwedge \{x \in \mathbb{R} \mid \forall q \in g_\epsilon^{\partial S}(p) : f_v(q) > x\}$$

For the $(d-1)$ -plane $H = \{q \in \mathbb{R}^d \mid f_v(q) = a\}$, there is a bounded connected component $k \in \kappa(S^c \setminus H)$, implying that S is not a hull. We conclude that if S is a hull of R , then every strict local maximum of f_v for some $v \in \mathbb{S}^{d-1}$ over S^- is a strict local maximum of f_v over R^- , and there is no strict local maximum of f_v over S^{c-} .

We now show that if S is not a hull, then there is a point $p \in \partial S$ and direction $v \in \mathbb{S}^{d-1}$ such that either p is a strict local maximum of height function f_v over S^- but p is not a strict local maximum f_v over R^- , or p is a strict local maximum over S^{c-} . Suppose there exists a $(d-1)$ -plane H such that there is a connected component $k \in \kappa(S \cap (S \setminus H))$ where $k \cap R = \emptyset$. Then there exists a point $p \in (\partial \text{CH}(k) \cap \partial k) \setminus H$ that is a strict local maximum of f_v for some direction $v \in \mathbb{S}^{d-1}$ over S^- . This is because $\partial \text{CH}(k) \setminus H$ must include a locally convex point, and $\partial \text{CH}(k) \setminus \partial k$ is not convex. Although p is a strict local maximum over S^- , point p is not a strict local maximum of f_v over R^- because $k \cap R = \emptyset$. Similarly, if there exists a $(d-1)$ -plane H such that there is a connected component $k \in \kappa(S^c \cap (S^c \setminus H))$ where k is bounded, then there exists a point $p \in (\partial \text{CH}(k) \cap \partial k) \setminus H$ that is a strict local maximum of height function f_v for some direction $v \in \mathbb{S}^{d-1}$ over S^{c-} .

If S is not a hull of R , there exists a point p that is a strict local maximum height function f_v for some $v \in \mathbb{S}^{d-1}$ over S^- , although p is not a strict local maximum of

f_v over R^- , or there exists strict local maximum of f_v over S^{c-} . We conclude that if there is no point p and direction $v \in \mathbb{S}^{d-1}$ such that p is a strict local maximum of height function f_v over ∂S but p is not a strict local maximum of f_v over R^- , then S is a hull of R . \square

Remark 12.5. In the following definition, we only consider sets whose boundaries have finite $(d - 1)$ -dimensional Hausdorff measure. (The elements of $\mathbb{R}^>$ are finite, so $\infty \notin \mathbb{R}^>$.) We require that S be topologically regular (Definition 8.25) because sets with empty interiors can have ambiguous or nonexistent enclosing sets of minimal boundary measure. The topological regularity requirement also excludes sets with lower-dimensional “hairs” or “cracks” from consideration.

Our objective is to show that sets with diverse geometric properties are hulls under our definition. As we establish in Chapter 13, the geometric properties of minimal hulls differ substantially from those of convex hulls. Smooth minimal hulls have zero mean curvature away from their constraints, making them saddle-shaped, while we conjecture that tight hulls have zero Gaussian curvature away from their constraints, making them developable. Minimal hulls minimize boundary measure, while tight hulls minimize slack. Informally, the difference is one of minimizing spatial variation versus minimizing normal variation over a set’s boundary.

Definition 12.4. For a topologically regular set $S \subseteq \mathbb{R}^d$ such that $R \subseteq S$ and $H^{d-1}(\partial S) \in \mathbb{R}^>$, if there is no topologically regular set $T \subseteq \mathbb{R}^d$ such that $R \subseteq T$ and $H^{d-1}(\partial T) < H^{d-1}(\partial S)$, then S is a *minimal hull* of R .

Remark 12.6. The proof’s argument is that if we can shave a connected component from a hull $S \subseteq \mathbb{R}^d$ of $R \subseteq S$, then the flat patch where the plane trims the component has less boundary measure than the remainder of the trimmed component’s boundary, demonstrating that the hull’s boundary measure is not minimal.

Theorem 12.3. *If $S \subseteq \mathbb{R}^d$ is a minimal hull of $R \subseteq S$, then S is a hull of R .*

PROOF. Suppose that there is a $(d - 1)$ -plane $H \subseteq \mathbb{R}^d$ such that there is a connected component $k \in \kappa(S \cap (S \setminus H))$ where $k \cap R = \emptyset$. Then $H^{d-1}(k \cap H) < H^{d-1}(\partial k \setminus H)$, so $H^{d-1}(S \setminus k) < H^{d-1}(S)$, implying S does not minimize boundary measure over sets containing R . If instead there is bounded connected component $k \in \kappa(S^c \cap (S^c \setminus H))$, then $H^{d-1}(k \cap H) < H^{d-1}(\partial k \setminus H)$, so $H^{d-1}(S \cup k) < H^{d-1}(S)$, again implying that S does not minimize boundary measure. We conclude that if S is a minimal hull of R , then S is a hull of R . \square

Remark 12.7. We now introduce relative hulls. Like a hull, a relative hull $S \subseteq \mathbb{R}^d$ encloses a set $R \subseteq S$, but it also excludes a set G disjoint from R . The boundary of the relative hull of S wraps around both R and G , forming a separating surface. Like the boundary of a hull, the boundary of a relative hull has no extraneous bumps corresponding to local height function maxima. The main difference between the definition of a hull and the definition of a relative hull is that a component of S^c produced by cutting a hull with a plane must be unbounded, while a component of S^c produced by cutting a relative hull may be unbounded or intersect G . In many of the situations we discuss throughout this dissertation, the complement of G is bounded, in which case the definition becomes symmetric: every component of a hull S of R relative to G cut off by a plane intersects R , while every component of S^c cut off by a plane intersects G .

Definition 12.5. For $S \subseteq \mathbb{R}^d$ such that $R \subseteq S$, $G \subseteq S^c$, and $(\partial S)^\circ = \emptyset$, set S is a *hull of R relative to G* if and only if for every $(d - 1)$ -plane H , every connected component $k \in \kappa(S \setminus H)$ intersects R and every connected component $k \in \kappa(S^c \setminus H)$ intersects G or is unbounded.

A set has the *relative hull property* if and only if it is relative hull.

Remark 12.8. We now prove theorems for relative hulls analogous to Theorems 12.1, 12.2, and 12.3. We condense sections of the proofs for relative hulls that repeat material from the proofs for hulls.

Lemma 12.2. *For all $R, G \subseteq \mathbb{R}^d$ such that $R \cap G = \emptyset$, the interior of the boundary of the convex hull of R relative to G is empty: $(\partial\text{CH}(R | G))^\circ = \emptyset$.*

PROOF. The proof is closely resembles that of Lemma 12.1, a key point being that

$$\text{CH}(R | G) = \bigcup_{X \in \wp(R)} \text{CH}(X | G)$$

If $(\partial\text{CH}(R | G))^\circ = \emptyset$, choose $b \in \mathbb{B}^{(\partial\text{CH}(R|G))^\circ}$ such that

$$\{q \in \mathbb{R}^d \mid (q \in \text{CH}(R | G)) \wedge (\exists p \in b : \overline{pq} \cap G = \emptyset)\} \neq \emptyset$$

If $\text{CH}(b \cap R | G) \subset b$, then there is a $p \in \text{CH}(b \cap R | G)$ and a $q \in \text{CH}(R | G) \setminus \text{CH}(b \cap R | G)$ such that $\overline{pq} \cap G = \emptyset$ and $\overline{pq} \not\subseteq \text{CH}(R | G)$, implying $\text{CH}(R | G)$ is not convex relative to G . If $b = \text{CH}(b \cap R)$, then b lies in the interior of $\text{CH}(R | G)$ rather than in its boundary. \square

Remark 12.9. The proof that a relative convex hull is a relative hull closely follows the proof of Theorem 12.1 that establishes that convex hulls are hulls. We show that if shaving off a piece of a convex hull $S \subseteq \mathbb{R}^d$ of $R \subseteq S$ relative to $G \subseteq S^c$ makes it no longer relatively convex, then S was not convex relative to G before we shaved the piece, contradicting the proposition that S is a relative convex hull. If instead S remains convex relative to G after the piece is removed, then it is not a relative convex hull because it is not the smallest convex set relative to G containing R . If we can shave off a bounded connected component from S^c that is disjoint from G , then S is not convex relative to G because its does not contain line segments with endpoints in S that pass through the empty component. Consequently relative convex hulls are relative hulls.

Theorem 12.4. *For $R, G \subseteq \mathbb{R}^d$ such that $R \cap G = \emptyset$, the convex hull of R relative to G is a hull of R relative to G .*

PROOF. For all $R, G \subseteq \mathbb{R}^d$ such that $R \cap G = \emptyset$, $R \subseteq \text{CH}(R | G)$ because the relative convex hull of R relative to G is the convex hull of R in a convex structure. The convex hull of R relative is disjoint from G because by definition it is a union of line segments disjoint from G . By Lemma 12.2, $(\partial \text{CH}(R | G))^\circ = \emptyset$.

If there exists a $(d - 1)$ -plane such that a connected component $k \in \kappa(\text{CH}(R | G) \setminus H)$ does not intersect R , then $\text{CH}(R | G) \setminus k$ must be convex relative to G for reasons analogous to those given in the proof of Theorem 12.1. If $\text{CH}(R | G) \setminus k$ is not convex relative to G , there exist points $p, q \in \text{CH}(R | G) \setminus k$ such that $\overline{pq} \not\subseteq \text{CH}(R | G) \setminus k$, but $\overline{pq} \subseteq \text{CH}(R | G)$. Then $\overline{pq} \cap (\partial k / H) \neq \emptyset$, so $\overline{pq} \not\subseteq \text{CH}(R | G)$, contradicting that $\text{CH}(R | G)$ is convex relative to G . Then $\text{CH}(R | G) \setminus k$ is convex relative to G , but $\text{CH}(R | G) \setminus k \subset \text{CH}(R | G)$, contradicting that $\text{CH}(R | G)$ is the smallest convex set relative to G containing R when all such sets are ordered by set inclusion. We conclude that every connected component of $\text{CH}(R | G) \setminus H$ intersects R .

As in the proof of Theorem 12.1, if there exists a $(d - 1)$ -plane H such that a connected component $k \in \kappa(\text{CH}(S)^c \setminus H)$ is bounded and $k \cap G = \emptyset$, then there exist points $p, q \in \text{CH}(R | G)$ such that $\overline{pq} \cap k \neq \emptyset$. Because $\overline{pq} \cap k \neq \emptyset$, we have $\overline{pq} \not\subseteq \text{CH}(R | G)$, and because $p, q \in \text{CH}(R | G)$, we have $\overline{pq} \cap G = \emptyset$. Then $\text{CH}(R | G)$ is not convex relative to G , which is a contradiction. We conclude that every connected component of $\text{CH}(R | G) \setminus H$ intersects G or is unbounded. The convex hull of R relative to G is then a hull of R relative to G . \square

Remark 12.10. Below we prove the Relative Shaving Lemma. We rely on the reasoning from the proof of the Shaving Lemma that demonstrates that for a set $S \subseteq \mathbb{R}^d$ with $R \subseteq S$ and $G \subseteq S^c$, the existence of a strict height function maximum on ∂S that is not also a strict height function maximum on R^- or G^- implies that we can shave a bounded component of S or S^c that is disjoint from R and G . As this proves that S is not a hull of R relative to G , a hull S of R relative to G has no height maximum on ∂S that is not also a height function maximum on R^- or G^- .

If S is not a hull of R relative to G , then there is a bounded component of S or S^c disjoint from R and G that we can shave with a cutting plane. We show that a height function maximum on ∂S exists on such a component that is not a height function maximum on R^- or G^- . If there is no strict height function maximum on ∂S that is not also a strict height function maximum on R^- or G^- , then S is a hull of R relative to G .

Theorem 12.5. (*Relative Shaving Lemma*) For $S \subseteq \mathbb{R}^d$ such that $R \subseteq S$, $S \subseteq G^c$, and $(\partial S^\circ) = \emptyset$, the set S is a hull of R relative to G if and only if for every point $p \in \partial S$ and direction $v \in \mathbb{S}^{d-1}$ such that p is a strict local maximum of height function f_v on S^- , point p is a strict maximum of f_v on R^- , while if p is a strict local maximum of f_v on S^{c-} , then p is strict local maximum on G^- .

PROOF. Suppose that p is a strict local maximum of height function f_v on S^- for some $v \in \mathbb{S}^{d-1}$. but p is not a strict local maximum of f_v on R^- . By our reasoning in the proof of Theorem 12.2, there exists a $(d-1)$ -plane H such that there is a bounded connected component $k \in \kappa(S \setminus H)$ with $k \cap R = \emptyset$. Then S is not a hull of R relative to G because k is a connected component of $S \setminus H$ disjoint from R .

Instead suppose p is a strict local maximum of f_v on S^{c-} , but p is not a strict local maximum of f_v on G^- . Then S is not a hull of R relative to G because there exists a $(d-1)$ -plane H such that there is a bounded connected component $k \in (S^c \setminus H)$ disjoint from G . We conclude that if S is a hull of R relative to G , every strict local height function maximum on ∂S is a strict local height function maximum on R^- or G^- .

We now show that if S is not a hull of R relative to G , there is a point $p \in \partial S$ and direction $v \in \mathbb{S}^{d-1}$ such that p is a strict local maximum of height function f_v on ∂S but p is not a strict local maximum of f_v on R^- or G^- . Suppose that there exists a $(d-1)$ -plane H such that there is a connected component $k \in \kappa(S \cap (S \setminus H))$ where $k \cap R = \emptyset$. As in our proof of Theorem 12.2, there is a point $p \in (\partial \text{CH}(k) \cap \partial k) \setminus H$

that is a strict local maximum of a height function f_v for some $v \in \mathbb{S}^{d-1}$ over S^- . However, point p is not a strict local maximum of R^- because $k \cap R = \emptyset$, and p is not a strict local maximum on G^- because $k \subseteq S$ and $G \subseteq S^c$, so $k \cap G = \emptyset$.

Now suppose there exists a $(d-1)$ -plane H such that there is a connected component $k \in \kappa(S^c \cap (S^c \setminus H))$ where k is bounded and $k \cap G = \emptyset$. Then there exists a point $p \in (\partial\text{CH}(k) \cap \partial k) \setminus H$ that is a strict local maximum of height function f_v for some $v \in \mathbb{S}^{d-1}$ over S^c . However, point p is not a strict local maximum of f_v over G^- because $k \cap G = \emptyset$, and p is not a strict local maximum on R^- because $k \subseteq S^c$ and $R \subseteq S$.

If S is not a hull of R relative to G , there exists a point $p \in \partial S$ and direction $v \in \mathbb{S}^{d-1}$ such that p is a strict local maximum of height function f_v over ∂S but p is not a strict local maximum of f_v over R^- or G^- . We conclude that if there is no point $p \in \partial S$ and direction $v \in \mathbb{S}^{d-1}$ such that p is a strict local maximum of height function f_v over ∂S but not a strict local maximum of f_v over R^- or G^- , then S is a hull of R relative to G . \square

Remark 12.11. We now define the relative version of a minimal hull. In our earlier work on tightening [68], we proposed using relative minimal hulls to generalize tightening to three and higher dimensions. That choice was inspired by the fact that relative convex hulls have locally minimal boundary measure in two dimensions. In higher dimensions, however, relative convex hulls do not minimize boundary measure and minimal hulls do not maximize convexity. This motivates our definition of tight hull in Chapter 13. Like relative convex hulls, tight hulls are maximally convex, and like minimal hulls they behave symmetrically with respect to set complement. The symmetry of minimal hulls derives from the fact that they minimize boundary measure, which is invariant under set complement.

As in the proof of Theorem 12.3, we argue that if we can shave a bounded connected component from a minimal hull $S \subseteq \mathbb{R}^d$ relative to $R \subseteq S$ and $G \subseteq S^c$, then the flat patch left by the cutting plane has a smaller boundary measure than the

remainder of the boundary of the shaved component. If S is not a relative hull, then S is not a relative minimal hull, so if S is a relative minimal hull, S is a relative hull.

Definition 12.6. For a topologically regular set $S \subseteq \mathbb{R}^d$ such that $R \subseteq S$, $S \subseteq G^c$, and $H^{d-1}(\partial S) \in \mathbb{R}^>$, if there is no topologically regular set $T \subseteq \mathbb{R}^d$ such that $R \subseteq T$, $T \subseteq G^c$, and $H^{d-1}(\partial T) < H^{d-1}(\partial S)$, then S is a *minimal hull of R relative to G* .

Theorem 12.6. *If $S \subseteq \mathbb{R}^d$ is a minimal hull of $R \subseteq S$ relative to $G \subseteq S^c$, then S is a hull of R relative to G .*

PROOF. Suppose there is a $(d-1)$ -plane H such that there is a connected component $k \in \kappa(S \cap (S \setminus H))$ where $k \cap R = \emptyset$; note that because $S \cap G = \emptyset$, we have $k \cap G = \emptyset$ as well. Then $H^{d-1}(k \cap H) < H^{d-1}(\partial k \setminus H)$, so $H^{d-1}(S \setminus k) < H^{d-1}(S)$, implying S does not minimize boundary measure over hulls of R relative to G . Suppose instead there is bounded connected component $k \in \kappa(S^c \cap (S^c \setminus H))$ where $k \cap G = \emptyset$; because $S^c \cap R = \emptyset$, we have $k \cap R = \emptyset$ as well. Then $H^{d-1}(k \cap H) < H^{d-1}(\partial k \setminus H)$, so $H^{d-1}(S \cup k) < H^{d-1}(S)$, implying S does not minimize boundary measure. We conclude that if S minimizes boundary measure over all sets that contain R and exclude G , then S is a hull of R relative to G . \square

Remark 12.12. In the next section we shift from characterizing hull boundaries with the Shaving Lemma and Relative Shaving Lemma to instead using convexity normals and support. Rather than requiring a local height function maximum of a hull to coincide with a local height function maximum of R^- or G^- , we require that locally convex points on a hull be “fully supported” by convexity normals from R and G . While this approach has a weaker relationship to the two-piece property than that provided by the Shaving Lemma, it allows us to quantify the extent to which R and G support the boundary of a hull, envisioned as a membrane stretched between them. That later proves central in our definition of tight hulls.

12.2. Support and convexity normals

Remark 12.13. We begin by defining local convexity and strict convexity. A locally convex point is a height function maximum over a small geodesic neighborhood, while a strictly convex point is instead a strict height function maximum. Not every locally convex point is strictly convex, but every strictly convex point is a locally convex.

Definition 12.7. For a set $S \subseteq \mathbb{R}^d$, a point $p \in \partial S$ is *locally convex* if and only if there exists a $v \in \mathbb{S}^{d-1}$ such that p is a local maximum of height function f_v over S^- . Point p is *locally concave* point of S if and only if it is a locally convex point of S^{c-} .

Example 12.1. A point on the boundary of an infinite solid right circular cylinder in \mathbb{R}^3 is locally convex. A point on a halfspace boundary is both locally convex and locally concave.

Definition 12.8. For a set $S \subseteq \mathbb{R}^d$, a point $p \in \partial S$ is *strictly convex* if and only if there exists a $v \in \mathbb{S}^{d-1}$ such that p is a strict local maximum of height function f_v over S^- . Point p is *strictly concave* point of S if and only if it is a strictly convex point of S^{c-} .

Example 12.2. A point on the boundary of a solid ball is strictly convex. A point on the boundary of an infinite solid right circular cylinder is not strictly convex, nor is a point on a halfspace boundary.

Remark 12.14. The definitions of local and strict convexity above focus on convexity at a point rather than over the neighborhood of a point. A polyhedron, for instance, may have a convex vertex that has both convex and concave incident edges. Although the vertex is convex, its neighborhood is not.

Nevertheless, if every point on a topological regular set (Definition 8.25) — informally, a solid — is locally convex, then we can conclude that the closure of each of its connected components is convex. The theorem is not true for arbitrary sets. If we draw a nonconvex curve $X \subseteq \mathbb{R}^3$ on an infinite right circular cylinder $C \subseteq \mathbb{R}^3$, so

$X \subseteq C$, every point $p \in X$ is a local maximum of the height function defined by the vector pointing to p from the closest point to p on the axis of C . Then set X is not convex despite the fact that each of its points is locally convex.

Theorem 12.7. *For a topologically regular set $S \subseteq \mathbb{R}^d$, if every $p \in \partial S$ is locally convex, then for each $k \in \kappa(S)$, the component k^- is convex.*

PROOF. If k^- is not convex, then there exist $p, q \in k^-$ such that $\overline{pq} \not\subseteq k^-$. There exists a minimum-length path $x \in k^-$ connecting p and q . There also exists a $w \in x \cap (\partial k \setminus p \setminus q)$ that is not a locally convex point of k^- . We have $x \neq \overline{pq}$ because $x \subseteq k^-$ and $\overline{pq} \not\subseteq k^-$. Consequently, x is not a line segment. A point of x that is not in the interior of a line segment coincides with a point on ∂k that is not locally convex. Otherwise, it would be possible to locally reduce the length of x by deforming it within k^- . Because our inference that a point on ∂k is not locally convex contradicts the hypothesis that every point on ∂S is locally convex, we conclude that the closure of every connected component of S is convex. \square

Remark 12.15. For a set $S \subseteq \mathbb{R}^d$, we introduce a definition of the convexity of a normal $v \in N_p^S$ at a point $p \in \partial S$. The convexity of v at p is the fraction of 2-dimensional sections through the line $\overleftrightarrow{p(p+v)}$ in which S is strictly convex at p . Qualitatively, convexity measures the extent to which a point behaves like a bump along a given normal direction.

Definition 12.9. For a set $S \subseteq \mathbb{R}^d$, let $v \in \mathbb{S}^{d-1}$ be a normal to S at $p \in \partial S$. Let X be the set of all 2-dimensional affine subspaces ϖ of \mathbb{R}^d containing the line $\overleftrightarrow{p(p+v)}$ for which $\Psi(p)$ is a strictly convex point of $\Psi(S)$ under an isometric mapping Ψ of ϖ to \mathbb{R}^2 . Then the *convexity* of v at p for S is defined as

$$\frac{H^{d-2}(X)}{H^{d-2}(\mathbb{S}^{d-2})}$$

For a set $S \subseteq \mathbb{R}^d$ and point $p \in \partial S$, a *convexity normal* at p to S is a normal $v \subseteq N_p^S$ scaled by the convexity of v at p for S .

Notation 18. For a set $S \subseteq \mathbb{R}^d$ and point $p \in \partial S$, let CV_p^S denote the set of convexity normals to S at p .

Example 12.3. For $S \subseteq \mathbb{R}^d$, the convexity of every normal at strictly convex point $p \in \partial S$ is one. The convexity of the unique normal at each point on the inner wall of a round torus is less than one. Suppose a and b are two infinite solid right circular cylinders of equal radius whose axes intersect at a right angle. The convexity of every normal along their intersection curve is less than one, but different normals at the same point on the intersection curve can have different convexities.

Remark 12.16. We define the concept of support to capture the idea of the boundary of one set being stretched over another set, which provides it more or less support depending on the convexity of its normals. We show that the convex hull of a set $S \subseteq \mathbb{R}^d$ is supported by S . In two dimensions, we show that a set $S \subseteq \mathbb{R}^d$ is a hull of $R \subseteq S$ relative to $G \subseteq S^c$ if and only if S is supported by R and G . This formalizes the perception that the boundary components of a relative hull in two dimensions behave like rubber bands in tension, stretched over the convex points of R and G .

Definition 12.10. For a set $S \subseteq \mathbb{R}^d$, a point $p \in \partial S$, and convexity normal $v \in CV_p^S$, and $T \subseteq S$, the *support* of v from T is defined as $\bigvee_{x \in CV_p^T} (x \cdot v)$. Point $p \in \partial S$ is unsupported if and only if its support is zero, and $X \subseteq \partial S$ is unsupported if and only if every $p \in X$ is unsupported.

Notation 19. For $S \subseteq \mathbb{R}^d$, $T \subseteq S$, $p \in \partial S$, and $v \in N_p^S$, let $SP_p^{S,T}(v)$ denote the support of v at p from T .

Definition 12.11. [### normal field reachability] The *weighted support* of a normal $v \in \mathbb{S}^{d-1}$ to $S \subseteq \mathbb{R}^d$ over $X \subseteq \partial S$ from $T \subseteq S$, denoted $WS_X^{S,T}(v)$, is defined as

$$WS_X^{S,T}(v) = \sum_{k \in NP^S(v)} \left[\bigvee \{ \{ NC_p^S(v) SP_p^{S,T}(v) \mid p \in k \cap X \} \cup \{0\} \} \right]$$

The *supported slack* of S from T over X , denoted $\text{SS}_X^{S,T}$, is defined as

$$\text{SS}_X^{S,T} := \int_{v \in \mathbb{S}^{d-1}} \text{WS}_X^{S,T}(v)$$

For a set $S \subseteq \mathbb{R}^d$ with $R \subseteq S$ and $G \subseteq S^c$, the *supported slack* of S from R and G is the sum of the supported slack of S over ∂S from R and the supported slack of S^c over ∂S from G . The *unsupported slack* of S is the difference between the slack of S and the supported slack of S .

Theorem 12.8. *The supported slack of $S \subseteq \mathbb{R}^d$ over $X \subseteq \partial S$ from $T \subseteq S$ is less than or equal to the slack of S over X .*

PROOF. The definitions of weighted support and normal count (Definition 10.13) are identical except for the factor $\text{SP}_p^{S,T}(v)$ in the equation for weighted support. A vector's support is always in $[0, 1]$, so $\text{SP}_p^{S,T}(v) \in [0, 1]$. The definitions of supported slack and slack (Definition 10.14) are also almost identical. We compute supported slack by integrating weighted support over the sphere of directions, while we compute slack by integrating normal count. Because weighted support is less than or equal to normal count, supported slack is less than or equal to slack. \square

Corollary 12.1. *For $S \subseteq \mathbb{R}^d$ with $R \subseteq S$ and $G \subseteq S^c$, the supported slack of S from R and G is less than or equal to the slack of S .*

PROOF. The normal count of a vector $v \in \mathbb{S}^{d-1}$ for $S \subseteq \mathbb{R}^d$ at $p \in \partial S$ is equal to the normal count of $-v$ for S^c at p : $\text{NC}_p^S(v) = \text{NC}_p^{S^c}(-v)$. In addition, the sum of the support of a normal $v \in \mathbb{S}^{d-1}$ at $p \in \partial S$ from $R \subseteq S$ and the support of $-v$ at p from $G \subseteq S^c$ is less than or equal to one. This is because the convexity of v is the fraction of sections in which S is strictly convex, while the convexity of $-v$ is the fraction of sections in which S is strictly concave. If the sum of the convexities of v and $-v$ were greater than one, it would imply the existence of a section that was both strictly convex and strictly concave, a contradiction. The sum of the support

of v from R and the support of $-v$ from G is less than or equal to the sum of the convexities of v and $-v$, and so it is also less than or equal to one. Then

$$\text{WS}_p^{S,R}(v) + \text{WS}_p^{S^c,G}(-v) \leq \text{NC}_p^S(v)$$

The supported slack of S from R and G is the integral over all $v \in \mathbb{S}^{d-1}$ of the left side of the inequality, while the slack of S is the integral over all directions of the right side of the inequality. We conclude that the supported slack of S from R and G is less than or equal to the slack of S . \square

Lemma 12.3. *For $S \subseteq \mathbb{R}^d$, if $T \subseteq S$ and $p \in \partial S \cap \partial T$, then $N_p^S \subseteq N_p^T$.*

PROOF. ### \square

Lemma 12.4. *For a closed set $S \subseteq \mathbb{R}^d$, if $X \subseteq \partial S$ is locally but not strictly convex for all $p \in X$, then the slack of S over X is zero.*

PROOF. ### \square

Theorem 12.9. *For all $S \subseteq \mathbb{R}^d$, the convex hull of S is supported by S .*

PROOF. At a strictly convex point $p \in \text{CH}(S)$, the convexity of every $v \in \text{CV}_p^{\text{CH}(S)}$ is one. The convexity at p for S is also one, because if there were a section through $\overleftrightarrow{p(p+v)}$ that were not strictly convex it would imply that $S \not\subseteq \text{CH}(S)$, a contradiction. Therefore the support of v from S is one. Then the weighted support of $\text{CH}(S)$ from S over the strictly convex subset of $\partial \text{CH}(S)$ is equal to the slack of $\text{CH}(S)$ over the strictly convex subset of $\partial \text{CH}(S)$. By Lemma , the slack over the subset of $\partial \text{CH}(S)$ that is locally but not strictly convex is zero, so the slack of $\text{CH}(S)$ is equal to the slack of the strictly convex subset of $\partial \text{CH}(S)$. We conclude that the convex hull of S is supported by S . \square

Theorem 12.10. *If $T \subseteq \mathbb{R}^d$ is compact and supported by $S \subseteq T$, then $T = \text{CH}(S)^-$*

PROOF. ###

□

Lemma 12.5. *If $S \subseteq \mathbb{R}^2$ is a hull of R relative to G , the unsupported subset of ∂S consists of line segments.*

PROOF. If some connected subset $X \subseteq \partial S$ is not a line segment, then there exists a point $p \in X$ that is a strict maximum on S or S^c of height function f_v for some $v \in \mathbb{S}$. If p is a strict maximum on S , then p does not lie on ∂R , because that would contradict our hypothesis that X is unsupported. Consequently, there exists a line H such that there is a connected component $k \in \kappa(S \cap (S \setminus H))$ with $p \in \partial k$ such that $k \cap R = \emptyset$, contradicting our hypothesis that S is a hull of R relative to G .

Similarly, if p is a strict maximum on S^c , then p does not lie on ∂G , and there is a line H and component $k \in \kappa(S^c \cap (S^c \setminus H))$ with $p \in \partial k$ such that $k \cap G = \emptyset$. This again contradicts our hypothesis that S is a hull of R relative to G , so we conclude that every connected subset of the unsupported subset of ∂S is a line segment. □

Theorem 12.11. *A set $S \subseteq \mathbb{R}^2$ with $R \subseteq S$, $S \subseteq G^c$, and $(\partial S)^\circ = \emptyset$ is a hull of R relative to G if and only if S is supported by R and G .*

PROOF. If $S \subseteq \mathbb{R}^2$ is a hull of $R \subseteq S$ relative to $G \subseteq S^c$, then by the Relative Shaving Lemma and Definition 12.7, every strictly convex point $p \in \partial S$ is a strictly convex point on the boundary of R . For any $v \in N_p^S$ there is one 2-dimensional affine subspace of \mathbb{R}^2 containing the line $\overleftrightarrow{p(p+v)}$, namely \mathbb{R}^2 , so the convexity of every normal to S at p is one. By the same argument, the convexity of every normal to R at p is one. By Lemma 12.3, we have $N_p^S \subseteq N_p^R$, so the supported slack of the strictly convex subset of ∂S is equal to its slack.

Similarly, every strictly concave point $p \in \partial S$ is a strictly convex point on the boundaries of S^c and G . The convexity of every normal to S^c at p is one, as is the convexity of every normal to G at p . We have $N_p^{S^c} \subseteq N_p^G$, so the supported slack of the strictly concave subset of ∂S is equal to its slack.

The remainder of ∂S is unsupported, because if $p \in \partial S$ is not a strictly convex point of S , the convexity normal to S at p is zero. Similarly, if p is not a strictly convex point of S^c , the convexity normal to S^c at p is zero. By Lemma 12.5, unsupported subset of ∂S consists of line segments, so its slack is zero. We conclude that if S is a hull of R relative to G , it is supported by R and G .

Now suppose S is supported by R and G . Assuming S is not a hull of R relative to G , then there exists a $(d - 1)$ -plane H such that exists a connected component $a \in \kappa(S \cap (S \setminus H))$ such that $a \cap R = \emptyset$ or there exists a bounded connected component $b \in \kappa(S^c \cap (S^c/H))$ such that $b \cap G = \emptyset$. If component a exists, the slack over $\partial a \setminus H$ is nonzero but the supported slack over ∂a from R is zero, contradicting our hypothesis that S is supported by R and G . Similarly, if component b exists the slack over $\partial b \setminus H$ is nonzero but the supported slack of ∂b from G is zero. We conclude that if S is supported by R and G then S is a hull of R relative to G . \square

Definition 12.12. A k -dimensional affine subspace $H \subseteq \mathbb{R}^d$ for $0 < k < d$ is tangent to $S \subseteq \mathbb{R}^d$ at $p \in \partial S$ if and only if $N_p^H \cap N_p^S \neq \emptyset$.

Corollary 12.2. *If $S \subseteq \mathbb{R}^2$ is a hull of R relative to G and for $p, q \in \partial R \cup \partial G$, segment $\overline{pq} \subseteq \partial S$ is unsupported, then \overline{pq} is tangent to $\partial R \cup \partial G$ at p and q .*

PROOF. Suppose $v \in \mathbb{S}$ is the normal to S along \overline{pq} , and $p \in \partial R$. If \overline{pq} is not tangent to ∂R or ∂G at p , then ∂S has an unsupported fan of normals interpolating between the normal $v \in \mathbb{S}$ to S along \overline{pq} and N_p^R . This implies that S is not supported by R and G . Then S is not a hull of R relative to G , contradicting our hypothesis. The argument for $p \in \partial G$ is similar, and p and q are interchangeable in either case. We conclude that every unsupported segment of the boundary of a hull of R relative to G is tangent to $\partial R \cup \partial G$ at its endpoints. \square

Corollary 12.3. *For $S \subseteq \mathbb{R}^2$ with $R \subseteq S$ and $(\partial S)^\circ = \emptyset$, then S is a hull of R if and only if S is supported by R .*

PROOF. This is a special case of Theorem 12.11 with $G = \emptyset$. □

Corollary 12.4. *For compact $S \subseteq \mathbb{R}^2$ with $R \subseteq S$, if S is a hull of R then $S = \text{CH}(R)^-$.*

PROOF. This is a consequence of Theorem 12.10 and Corollary 12.3. □

Theorem 12.12. *A set $S \subseteq \mathbb{R}^d$ with $R \subseteq S$, $G \subseteq S^c$, and $(\partial S)^\circ = \emptyset$ is a hull of R relative to G if and only if there is no point $p \in \partial S$ such that p is a strictly convex point of S unsupported by R or p is a strictly concave point of S unsupported by G .*

PROOF. This theorem reformulates the Relative Shaving Lemma (Theorem 12.5,) so we only sketch its proof. A strictly convex point unsupported by R or a strictly concave point unsupported by G can be shaved off by a cutting plane, so a relative hull has no unsupported strictly convex or concave points. If it is possible to shave a piece from S with a cutting plane, then S has an unsupported strictly convex or concave point. Conversely, if S has no unsupported strictly convex or concave points, it is impossible to shave a piece from S , so S is a hull of R relative to G . □

Corollary 12.5. *Suppose $S \subseteq \mathbb{R}^3$ with $R \subseteq S$, $G \subseteq S^c$, and $(\partial S)^\circ = \emptyset$, and assume the normals to ∂S are differentiable. Then S is a hull of R relative to G if and only if the Gaussian curvature of the unsupported subset of ∂S is nonpositive.*

PROOF. Under the assumption that the normals to $\partial S \subseteq \mathbb{R}^3$ are differentiable, the locally convex and locally concave points on the boundary of S are precisely those with strictly positive Gaussian curvature. □

12.3. Topological and geometric distinguishability

12.4. Covers locally indistinguishible from hulls

—OVERVIEW

—ABSTRACT HULL GEOMETRY

We show that an abstract hull's locally convex boundary lies on the locally convex part of the boundary of R , while the hull's locally concave boundary lies on the locally convex part of the boundary of G . We refer to the remainder of the hull's boundary as unsupported, a term we define in Chapter 13. Because it is neither locally convex nor locally concave, an abstract hull's unsupported boundary is either saddle-shaped or developable.

We show that the unsupported boundary of a two-dimensional abstract hull consists of line segments tangent to locally convex boundary intervals from R and G .

—TOPOLOGICAL AND GEOMETRIC DISTINGUISHIBILITY

This suggests that two-dimensional abstract hulls have similar local geometries, so we define locally indistinguishable sets and prove that two-dimensional abstract hulls are locally indistinguishable.

We define two abstract hulls of R relative to G as topologically distinct if we cannot deform the boundary of one into the other while holding R and G fixed during the deformation. We show that the number of topologically distinct abstract hulls of R relative to G is not bounded by a polynomial function of the number of connected components of R and G .

In two dimensions, we identify the relative convex hull as the union of all abstract hulls, while a tight hull is an abstract hull that minimizes slack. Although two-dimensional tight and relative convex hulls are locally indistinguishable, they may have distinct topologies. In higher dimensions, tight and relative convex hulls differ in both their local geometry and their topology, which motivates our definition of a cover.

—COVERS

Given a definition of a particular abstract hull, the corresponding set of covers contains all sets that are locally indistinguishable from the hull. In two dimensions, any two abstract hulls are locally indistinguishable, so the set of tight covers and the set of relative convex covers are both equal to the set of all abstract hulls.

In three dimensions, however, the set of tight covers and the set of relative convex covers are equal if and only if all covers in both sets are locally convex.

A hull's covers factor topology from its definition, retaining its local geometric properties. When a hull's geometric properties are more relevant to an application than its topology, its covers provide a topological choice that we may use to optimize the output's quality or facilitate its computation.

PART 3

TIGHT HULLS AND TIGHTENING

CHAPTER 13

TIGHT HULLS

CHAPTER 14

TIGHTENING

CHAPTER 15

THE MEDIAL COVER

CHAPTER 16

IMPLEMENTATION

PART 4

PRIOR ART, FUTURE WORK, CONTRIBUTIONS, AND CONCLUSION

CHAPTER 17

PRIOR ART: FAIRING, BLURRING, AND BLENDING

CHAPTER 18

FUTURE WORK: CONSTRUCTING 3D TIGHT HULLS

CHAPTER 19

CONTRIBUTIONS

CHAPTER 20

CONCLUSION

REFERENCES

- [1] ATTALI, D. 1997. r -Regular shape reconstruction from unorganized points. *Proceedings of the 13th Annual Symposium on Computational Geometry*, Nice, France, June 1997, J.-D. BOISSONNAT, Ed. ACM, New York, NY, 248-253.
- [2] BANCHOFF, T. 1971. The two-piece property and tight n -manifolds-with-boundary in E^n . *Transactions of the American Mathematical Society* 161, 259-267.
- [3] BANCHOFF, T. 1997. Tight submanifolds, smooth and polyhedral. In *Mathematical Sciences Research Institute Publications Volume 32: Tight and Taut Submanifolds*, T. CECIL AND S.-S. CHERN, Eds. Cambridge University Press, New York, NY.
- [4] BERTRAND, J., DIQUET, C.-F., AND PUISEUX, V. 1848. Démonstration d'un théorème de Gauss. *Journal de Mathématiques Pures et Appliquées* 13, 80-90.
- [5] BLUM, H. 1967. A transformation for extracting new descriptors of shape. In *Models for the Perception of Speech and Visual Form*, W. DUNN, Ed. MIT Press, Cambridge, MA, 362-380.
- [6] BOBENKO, A., SCHRÖDER, P., SULLIVAN, J., AND ZIEGLER, G., Eds. 2008. *Discrete Differential Geometry*. Birkhäuser Verlag, Boston, MA.
- [7] BREHM, U. AND KÜHNEL, W. 1982. Smooth approximation of polyhedral surfaces regarding curvatures. *Geometriae Dedicata* 12, 61-85.
- [8] CERVONE, D. 1997. Tightness for smooth and polyhedral immersions of the projective plane with one handle. In *Mathematical Sciences Research Institutes Publications Volume 32: Tight and Taut Submanifolds*, T. CECIL AND S.-S. CHERN, Eds. Cambridge University Press, New York, NY.
- [9] *CGAL: Computational Geometry Algorithms Library*, <http://www.cgal.org>.
- [10] CHAZAL, F., CONHEN-STEINER, D., LIEUTIER, A., AND THIBERT, B. 2007. Shape smoothing using double offsets. *Proceedings of the 2007 ACM Symposium on Solid and Physical Modeling*, Beijing, China, June 2007, ACM Press, New York, NY, 183-192.
- [11] CHAZELLE, B. 1982. A theorem on polygon cutting with applications. *Proceedings of the 23rd Annual Symposium on Foundations of Computer Science*, Chicago, IL, November 1982, IEEE Computer Society, New York, NY, 339-349.
- [12] CHAZELLE, B. 1984. Convex partitions of polyhedra: A lower bound and worst-case optimal algorithm. *SIAM Journal of Computing* 13, 488-507.

- [13] CHERN, S.-S. AND LASHOF, R. 1957. On the total curvature of immersed manifolds. *American Journal of Mathematics* 79, 306-318.
- [14] CHEW, L. 1987. Constrained Delaunay triangulations. *Proceedings of the 3rd Annual Symposium on Computational Geometry*, Waterloo, Canada, June 1987, ACM Press, 215-222.
- [15] CHOI, H., CHOI, S., AND MOON, H. 1997. Mathematical theory of medial axis transform. *Pacific Journal of Mathematics* 181, 57-88.
- [16] DESBRUN, M., MEYER, M., SCHRÖDER, P., AND BARR, A. 1999. Implicit fairing of irregular meshes using diffusion and curvature flow. *Proceedings of the 26th Annual Conference on Computer Graphics: SIGGRAPH 1999*, 317-324.
- [17] DESCARTES, R. ca. 1620. *Elementary Treatise on Polyhedra*. Unpublished.
- [18] DEY, T. 2007. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge University Press, New York, NY.
- [19] EDELSBRUNNER, H., KIRKPATRICK, D., AND SEIDEL, R. 1983. On the shape of a set of points in the plane. *IEEE Transactions on Information Theory IT-29*, 551-559.
- [20] EDELSBRUNNER, H., AND MUCKE, E. 1994. Three-dimensional alpha shapes. *ACM Transactions on Graphics* 13, 43-72.
- [21] EVANS, L. AND SPRUCK, J. 1991. Motion of level sets by mean curvature I. *Journal of Differential Geometry* 33, 635-681.
- [22] EWING, G. 1985. *Calculus of Variations with Applications*. Dover Publications, New York, NY.
- [23] FEDERER, H. 1969. *Geometric Measure Theory*. Springer-Verlag, Berlin, Germany.
- [24] FREY, W. 2002. Boundary triangulations approximating developable surfaces that interpolate a space curve. *International Journal of Foundations of Computer Science* 13, 285-302.
- [25] GORECKA-EGAN, E. Helen of Troy. In *Paper Sculpture*, by MCPHARLIN, P., 1944, Marquardt & Company, New York, NY, pp. 46.
- [26] GRAYSON, M. 1987. The heat equation shrinks embedded plane curves to round points. *Journal of Differential Geometry* 26, 285-314.
- [27] HAUSDORFF, F. 1918. Dimension und äusseres Mass. *Mathematische Annalen* 79: 157-179.
- [28] HEIJMANS, H. 1994. *Morphological Image Operators*. Academic Press, Boston, MA.
- [29] HERSHBERGER, J. AND SNOEYINK, J. 1994. Computing minimum length paths of a given homotopy class. *Computational Geometry: Theory and Applications* 4, 63-97.
- [30] JORDAN, C. 1892. Remarques sur les intégrales définies. *Journal de Mathématiques Pures et Appliquées* 8, 69-99.
- [31] KADUSHIN, R. 2006. Square dance coffee table. http://ronen-kadushin.com/Open_Design.asp.

- [32] KILIAN, M., FÖRY, S., CHEN, Z., MITRA, N., SHEFFER, A., AND POTTMANN, H. 2008. Curved folding. *ACM Transactions on Graphics 27: SIGGRAPH 2008*, 75.
- [33] KIM, B. AND ROSSIGNAC, J. 2005. GeoFilter: Geometric selection of mesh filter parameters. *Computer Graphics Forum 24: Eurographics 2005*, 295-302.
- [34] KLETTE, R. AND ROSENFELD, A. 2004. *Digital Geometry: Geometric Methods for Digital Picture Analysis*. Morgan Kaufmann, San Francisco, CA.
- [35] KÜHNEL, W. 2006. *Student Mathematical Library Volume 16: Differential Geometry: Curves — Surfaces — Manifolds*. American Mathematical Society, Providence, RI.
- [36] KUIPER, N. 1985. Minimal total curvature, a look at old and new results. *Russian Mathematical Surveys 40*, 49-55.
- [37] KUIPER, N. 1997. Geometry in curvature theory. In *Mathematical Sciences Research Institute Publications Volume 32: Tight and Taut Submanifolds*, T. CECIL AND S.-S. CHERN, Eds. Cambridge University Press, New York, NY.
- [38] LEE, D. 1978. *Proximity and Reachability in the Plane*. 1978. Doctoral dissertation, University of Illinois at Urbana-Champaign, Champaign, IL.
- [39] LEE, D. AND PREPARATA, F. 1984. Euclidean shortest paths in the presence of rectilinear barriers. *Networks 14*, 393-410.
- [40] LEVY, B. 2006. Laplace-Beltrami eigenfunctions: Towards an algorithm that “understands” geometry. *Proceedings of the IEEE International Conference on Shape Modeling and Applications: Shape Modeling International 2006*, Matsushima, Japan, June 2006, A BELYAEV, H. SUZUKI, AND M. SPAGNUOLO, Eds. IEEE Computer Society, Los Alamitos, CA, 13.
- [41] LIEUTIER, A. 2003. Any open bounded subset of \mathbb{R}^n has the same homotopy type than its medial axis. *Proceedings of the 8th Symposium on Solid Modeling and Applications*, Seattle, WA, June 2003, E. GERSHON AND V. SHAPIRO, Eds. The Association for Computing Machinery, New York, NY, 65-75.
- [42] MORETON, H. 1992. *Minimum Curvature Variation Curves, Networks, and Surfaces for Free-Form Shape Design*. Doctoral dissertation, University of California at Berkeley, Berkeley, CA.
- [43] MORETON, H. AND SEQUIN, C. 1992. Functional optimization for fair surface design. *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques: SIGGRAPH 1992*, 167-176.
- [44] ORZAN, A., BOUSSEAU, A., WINNEMÖLLER, H., BARLA, P., THOLLOT, J., AND SALESIN, D. 2008. Diffusion curves: A vector representation for smooth-shaded images. *ACM Transactions on Graphics 27: SIGGRAPH 2008*, 92.
- [45] PEANO, G. 1887. *Applicazioni Geometriche del Calcolo Infinitesimale*. Fratelli Bocca Editori, Turin, Italy.
- [46] PIZER, S., SIDDIQI, K., AND YUSHKEVICH, P. 2008. Introduction. In *Computational Imaging 37: Medial Representations: Mathematics, Algorithms, and Representations*, K. SIDDIQI AND S. PIZER, Eds. Springer, 1-34.

- [47] POTTMANN, H., ASPERL, A., HOFER, M., AND KILIAN, A. 2007. *Architectural Geometry*. Bentley Institute Press, Exston, PA.
- [48] POTTMANN, H., SCHIFTNER, P., BO, P., SCHMIEDHOFER, H., WANG, W., BALDASSINI, N., AND WALLNER, J. 2008. Freeform surfaces from single curved panels. *ACM Transactions on Graphics 27: SIGGRAPH 2008*, 76.
- [49] POTTMANN, H., WALLNER, J., HUANG, Q., AND YANG, Y. 2009. Integral invariants for robust geometry processing. *Computer Aided Geometric Design 26*, 37-60.
- [50] ROSE, K., SHEFFER, A., WITHER, J. CANI, M.-P., AND THIBERT, B. 2007. Developable surfaces from arbitrary sketched boundaries. *Proceedings of the 5th Eurographics Symposium on Geometry Processing*, Barcelona, Spain, July 2007, A. BELYAEV AND M. GARLAND, Eds. Eurographics Association, Aire-la-Ville, Switzerland, 163-172.
- [51] ROSE, N. 1985. *Mathematical Maxims and Minims*. Rome Press, Raleigh, NC.
- [52] ROSSIGNAC, J. 1985. *Blending and Offsetting Solid Models*. Doctoral dissertation, University of Rochester, Rochester, NY.
- [53] ROSSIGNAC, J. AND REQUICHA, A. 1984. Constant-radius blending in solid modeling. *ASME Computers in Mechanical Engineering 3*, 65-73.
- [54] SCARINCI, I. 2008. Financial scandal at the Guggenheim Bilbao. Translated by MAHABIR, A. *ARCADJA artMagazine*, <http://www.arcadja.com>.
- [55] SERRA, J. 1982. *Image Analysis and Mathematical Morphology, Volume 1*. Academic Press, New York, NY.
- [56] SETHIAN, J. 1999. *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science*. Cambridge University Press, New York, NY.
- [57] SKLANSKY, J. AND KIBLER, D. 1976. A theory of nonuniformly digitized pictures. *IEEE Transactions on Systems, Man, and Cybernetics 6*, 637-647.
- [58] SLOBODA, F. AND ZATKO, B. 2001. On approximation of Jordan surfaces in 3D. In *Lecture Notes in Computer Science Volume 2243: Digital and Image Geometry*, G. BERTRAND, A. IMIYA, AND R. KLETTE, Eds. Springer-Verlag, New York, NY, 365-386.
- [59] STRUIK, D. 1961. *Lectures on Classical Differential Geometry*. Dover Publications, New York, NY.
- [60] TAUBIN, G. 1995. Curve and surface smoothing without shrinkage. *Proceedings of the 5th Annual Conference on Computer Vision: ICCV 95*, 852.
- [61] TAUBIN, G. 1995. A signal processing approach to fair surface design. *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques: SIGGRAPH 1995*, 351-358.
- [62] TOUSSAINT, G. 1989. Computing geodesic properties inside a simple polygon. *Revue D'Intelligence Artificielle 3*, 9-42.

- [63] VAN DE VEL, M. 1993. *North-Holland Mathematical Library Volume 50: Theory of Convex Structures*. North-Holland, New York, NY.
- [64] WHITED, B., ROSSIGNAC, J., SLABAUGH, G., FANG, T., AND UNAL, G. 2007. Pearling: 3D interactive extraction of tubular structures from volumetric images. *MICCAI Workshop on Interaction in Medical Image Analysis and Visualization*, Brisbane, Australia, November 2007, A. PERSSON AND E. BENGTTSSON, Eds.
- [65] WHITED, B., ROSSIGNAC, J., SLABAUGH, G., FANG, T., AND UNAL, G. 2009. Pearling: Stroke segmentation with crusted pearl strings. *Journal of Pattern Recognition and Image Analysis* 19, 277-283.
- [66] WHITED, B., AND ROSSIGNAC, J. 2009. Relative blending. *Computer-Aided Design* 41, 456-462.
- [67] WILLIAMS, J. AND ROSSIGNAC, J. 2005. Mason: Morphological simplification. *Graphical Models* 67, 285-303.
- [68] WILLIAMS, J. AND ROSSIGNAC, J. 2007. Tightening: Morphological simplification. *International Journal of Computational Geometry and Applications* 17, 487-503.
- [69] WILLMORE, T. 1971. Tight immersions and total absolute curvature. *Bulletin of the London Mathematical Society* 3, 129-151.
- [70] WOLTER, F.-E. 1993. Cut locus and medial axis in global shape interrogation and representation. *Design Laboratory Memorandum 92-2*, Massachusetts Institute of Technology, Cambridge, MA.
- [71] WOLTER, F.-E. AND FRIESE, K.-I. 2000. Local and global geometric methods for the analysis, interrogation, reconstruction, modification and design of shape. *Proceedings of Computer Graphics International*, Geneva, Switzerland, June 2000, IEEE Computer Society, Los Alamitos, CA.
- [72] TRIM, P. Placeholder for a paper on variable-radius solid blending through medial axis trimming.