TIGHTENING AND BLENDING 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 2012

TIGHTENING AND BLENDING SUBJECT TO SET-THEORETIC CONSTRAINTS

Approved by:

Dr. Jarek Rossignac, Advisor
School of Interactive Computing
Georgia Institute of Technology

Dr. C. Karen Liu
School of Interactive Computing
Georgia Institute of Technology

Dr. Greg Turk
School of Interactive Computing
Georgia Institute of Technology

Dr. John McCuan
School of Mathematics
Georgia Institute of Technology

Dr. Helmut Pottmann
Geometric Modeling and Scientific Visualization Research Center
King Abdullah University of Science and Technology

Date Approved: May 14, 2012

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.

To my wife, a gift of paper.

ACKNOWLEDGEMENTS

I first thank my advisor, Jarek Rossignac, who has supplied me with intellectual support and guidance for several years. I would also like to acknowledge the students and visitors who have spent time in our MAGIC Lab, as well as the Geometry Group in the College of Computing at the Georgia Institute of Technology. I broadly thank my family, which has provided me with persistence, and in particular I express gratitude to my wife.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS v

LIST OF TABLES ix

LIST OF FIGURES x

NOTATION xv

SUMMARY xvii

PART 1. MASON AND TIGHTENING xviii

CHAPTER 1. INTRODUCTION 1

CHAPTER 2. DOCUMENT STRUCTURE 3

2.1. Division into parts 3

2.2. Part I 3

2.3. Part II 5

CHAPTER 3. MASON 7

3.1. Introduction 7

3.2. Mathematical morphology 8

3.3. Mason 12

3.4. Implementation 14

3.5. Discussion 23

3.6. Contributions 23

CHAPTER 4. TIGHTENING 25

4.1. Introduction 25

4.2. Definition and properties 25
4.3. Two-dimensional implementation 28
4.4. Three-dimensional implementation 31
4.5. Conclusion 33

CHAPTER 5. FAIRING, BLURRING, BLENDING 35
5.1. Introduction 35
5.2. Fairing 35
5.3. Blurring 37
5.4. Blending 40
5.5. Conclusion 45

PART 2. HULLS, BLENDS, AND COVERS 46

CHAPTER 6. TIGHT HULLS 47
6.1. Tight hulls generalize convex hulls 47
6.2. Comparison of hull definitions 50
6.3. Hulls of revolution 65
6.4. Symmetry with respect to set complement 71
6.5. The tablecloth problem 72

CHAPTER 7. TIGHT BLENDS 80
7.1. Tight blends bound curvature 80
7.2. Continuous tight blends of polygons 85
7.3. Tight blends simplify normal fields 87

CHAPTER 8. MEDIAL COVER 92

CHAPTER 9. IMPLEMENTATION 99

CHAPTER 10. APPLICATIONS 108
10.1. Convergent boundary estimation 108
10.2. Shape design 113

10.3. Blending 114
10.4. Normal field simplification 115
10.5. Further applications 116
CHAPTER 11. CONCLUSION 118
11.1. Introduction 118
11.2. Contributions 118
11.3. Discussion 120
11.4. Future work 121
Appendix A. Tight Hull Formalization 125
A.1. Polyhedra 125
A.2. Tightening 125
A.3. Slack 125
A.4. Convexity 126
A.5. Support 127
A.6. Tight hull 128
A.7. Tight hull properties 128
REFERENCES 130

LIST OF TABLES

1 Data from advantages of mason applied to a two-dimensional image. 17

2 Data from similar performances of operators applied to an input image. 17

3 Data from limited results obtained when applied to a two-dimensional image. 18

LIST OF FIGURES

3.1 Mason's advantages when applied to a two-dimensional input image. 16

3.2 Similar performance from operators applied to a two-dimensional image. 18

3.3 Limited results obtained from multiple operators applied to an image. 19

3.4 Mason resembles the input more closely than first opening or closing. 20

3.5 Insets illustrate opening and closing, with combinations and mason. 21

3.6 Nearly identical combinations of opening and closing. 22

3.7 Effects of progressively increasing the radius used in mason. 22

4.1 An r-tightening of 1102 by 832 pixels with a radius of 40 pixels. 30

4.2 An r-tightening of 1050 by 825 pixels with a radius of 75 pixels. 31

4.3 Tightening on a 128 x 128 x 279 voxel grid with a radius of 8 voxels. 32

4.4 Tightening on a 211 x 145 x 287 voxel grid with a radius of 12 voxels. 32

4.5 A sawtooth pattern from shrinking retained by opening and tightening. 33

6.1 A convex hull of points in the plane. 47

6.2 A hollow cylinder swept by rotating thin rectangles. 48

6.3 The convex hull of a hollow cylinder. 48

6.4 A separating surface with fluid-like behavior. 49

6.5 A separating surface like a membrane in tension. 49

6.6 Definition of a convex set. 51

6.7 Definition of a convex hull. 51

6.8 Convex hull construction by filtering subsets of Euclidean space. 52

6.9 Definition of relative convexity. 53

6.10 The convex hull of one simple polygon relative to another. 53

6.11 Embeddings of the torus that are and are not tight. 55

6.12 Embeddings of the sphere that are and are not tight. 56

6.13 Integral Gaussian curvature at a triangle mesh vertex. 56

6.14 Gaussian curvature from geodesic circles. 57

6.15 The total absolute curvature of a sphere and a cube are equal. 58

6.16 Sets that have and do not have the two-piece property. 58

6.17 A problem with defining tight hulls using the two-piece property. 59

6.18 Uncountably many sets may minimize constrained slack. 60

6.19 Unsupported slack for different candidate tight hulls. 61

6.20 Convex hulls have supported slack. 63

6.21 A singular geometric configuration yielding a nonunique tight hull. 64

6.22 An intersection of tight sets that is not tight. 64

6.23 The input sets R and G are symmetric about the same line. 66

6.24 The tight hull and convex hull of R relative to G are identical. 66

6.25 The sets R' and G' swept by rotating R and G. 67

6.26 The set R' consists of hollow and solid cylinders. 67

6.27 The set G' consists of a hollow cylinder and the complement of a cylinder. 68

6.28 The convex hull of R' relative to G'. 68

6.29 The tight hull of R' relative to G'. 69

6.30 Cross sections of the tight hull and convex hull of R' relative to G'. 70

6.31 The tight hull and convex hull of G' relative to R'. 71

6.32 The complements of the convex hull and tight hull of G' relative to R'. 72

6.33 The convex hull of four equal-radius balls. 73

6.34 The set R consists of four red balls, while G is one green ball. 73

6.35 The convex hull of R relative to G. 74

6.36 The tight hull of R relative to G. 74

6.37 Elliptic cylinders in the tight hull. 75

6.38 An incorrect tight hull of R relative to G with circular cylinders. 76

6.39 A spherical cap in the incorrect hull is larger than in the tight hull. 77

6.40 Singular arcs on the incorrect hull have excess normal variation. 78

6.41 The tight hull of a torus relative to a ball. 78

7.1 Hull boundaries press against locally convex constraint boundaries. 80

7.2 Planar hull boundaries are stable under constrained curvature flow. 81

7.3 The normals to a two-dimensional hull are supported. 82

7.4 The tight blend of a polygon. 83

7.5 The curvature of a hull's boundary is bounded by its constraints. 84

7.6 The normals to a tight blend may be continuous but nondifferentiable. 84

7.7 A tight blend of a polygon is a sequence of segments and arcs. 85

7.8 Arcs in a polygon's tight blend distribute the normal change at its vertices. 86

7.9 Increasing a tight blend radius slides arc endpoints together. 86

7.10 Changing the tight blend radius can change the result's topology. 87

7.11 The normal field to a polygon and its tight blend. 88

7.12 The Gauss offset of a nonconvex polygon. 89

7.13 The normal count from ray intersections with the Gauss offset. 89

7.14 Several paths connecting edge endpoints have identical slack. 90

7.15 Unsupported edges trim slack from incident arcs. 91

8.1 An algorithm for the relative convex hull of simple polygons. 93

8.2 Tight covers locally behave like tight hulls. 94

8.3 There may be an exponential number of tight covers. 94

8.4 Tight cover boundaries wind around holes in G^c. 95

8.5 Points added to R and G to define the medial cover. 96

8.6 The medial cover boundary is isotopic to loops equidistant from R and G. 97

8.7 The ordering of planar tight covers by set inclusion. 98

9.1 Red, green, and mixed triangles from the triangulation of $(R \cup G)^c$. 99

9.2 Mixed triangles form topological annuli. 100

9.3 The constrained Delaunay cover. 101

9.4 The constrained Delaunay cover depends on sampling. 102

9.5 Constrained Delaunay triangulation converges at high sampling rates. 103

9.6 Large triangle circumcenters converge to medial axis bifurcation points. 103

9.7 Insertion of bifurcation disk contact points as polygon vertices. 104

9.8 Bifurcation disk contact points make a cover invariant to sample insertion. 104

9.9 A medial cover algorithm can compute the convex hull. 105

9.10 Opening with an n-sided polygon has an output size linear in n. 106

9.11 Hausdorff error due to opening with a regular polygon rather than a disk. 106

9.12 Discrete analogs of bounded curvature in approximate blends. 107

10.1 Morphologically regular and irregular sets. 108

10.2 Red and green sets from samples. 109

10.3 Relative convex hull boundary measure converges with sampling density. 110

10.4 Peano-Jordan measure. 111

10.5 Midpoint reconstruction is nonconvergent. 112

10.6 Designs incorporating developable surfaces. 113

10.7 Developable surfaces interpolating boundary patches. 114

10.8 Small-radius tight blends provide unique normals. 115

10.9 Tight blends progressively simplify normal fields. 116

10.10 Repairing polygonal data with tight hulls. 117

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 set 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 $\mathbb{R}^{d+1}$
$\lvert\cdot\rvert$	absolute value
$\lVert\cdot\rVert$	Euclidean vector norm; for $v \in \mathbb{R}^d$, $\lVert v\rVert := \left(\sum_{i=1}^{d} {v_i}^2\right)^{\frac{1}{2}}$
$\bigvee$	supremum; least upper bound
$d(\cdot,\cdot)$	Euclidean distance function for points and point sets
$\mathrm{D}(\cdot)$	Hausdorff dimension
$\mathrm{H}^d(\cdot)$	d-dimensional Hausdorff measure
$O(f(x))$	set of functions asymptotically bounded from above by a multiple of $f(x)$
$\overline{pq}$	the closed line segment connecting points $p, q \in \mathbb{R}^d$
$\overrightarrow{pv}$	the ray with endpoint $p \in \mathbb{R}^d$ and direction $\frac{v}{\lVert v\rVert}$ 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; $\mathbb{B} := \left\{B_r(p) \mid (r \in \mathbb{R}^>) \wedge \left(p \in \mathbb{R}^d\right)\right\}$
$\mathbb{B}^S$	open balls contained in $S \subseteq \mathbb{R}^d$; $\mathbb{B}^S := \{b \in \mathbb{B} \mid b \subseteq S\}$
$\mathbb{B}(S)$	open balls with centers in $S \subseteq \mathbb{R}^d$; $\mathbb{B}(S) := \{b \in \mathbb{B} \mid \pi(b) \in S\}$
$\mathbb{B}_r$	open balls of radius $r \in \mathbb{R}^>$; $\mathbb{B}_r := \{b \in \mathbb{B} \mid \rho(b) = r\}$
$\mathbb{B}_r^S$, $\mathbb{B}_r(S)$	open balls of radius r in S, open balls of radius r with centers in S
$\mathbb{M}^S$	set of maximal balls in $S \subseteq \mathbb{R}^d$
MA^S	medial axis of $S \subseteq \mathbb{R}^d$; $\mathrm{MA}^S := \left\{\pi(b) \mid b \in \mathbb{M}^S\right\}$
S°	interior of $S \subseteq \mathbb{R}^d$; $S^\circ := \bigcup \mathbb{B}^S$
S^-	closure of $S \subseteq \mathbb{R}^d$; $S^- := S^{c\circ c}$
∂S	boundary of $S \subseteq \mathbb{R}^d$; $\partial S := S^- \cap S^{c-}$
$S\uparrow_r$	growing of $S \subseteq \mathbb{R}^d$ with radius $r \in \mathbb{R}^>$; $S\uparrow_r := \bigcup \mathbb{B}_r(S)$
$S\downarrow_r$	shrinking of $S \subseteq \mathbb{R}^d$ with radius $r \in \mathbb{R}^>$; $S\downarrow_r := S^c\uparrow_r{}^c$
$S\circ_r$	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 $r \in \mathbb{R}^>$; $S\bullet_r := S^c\circ_r{}^c$
$\mathcal{R}^S(p)$	regularity of $p \in S$ in $S \subseteq \mathbb{R}^d$; $\mathcal{R}^S(p) := \bigvee\{\{r \in \mathbb{R}^> \mid p \in S\circ_r{}^-\} \cup 0\}$
$\mathcal{M}_r S$	mortar of $S \subseteq \mathbb{R}^d$ for radius $r \in \mathbb{R}^>$; $\mathcal{M}_r S := S\bullet_r \cap S^c\bullet_r$
$\mathcal{T}_r S$	mason for $S \subseteq \mathbb{R}^d$ with radius $r \in \mathbb{R}^>$
$\mathrm{T}_r(S)$	unique tightening of $S \subseteq \mathbb{R}^d$ with radius $r \subseteq \mathbb{R}^d$
$\mathrm{AH}(\cdot)$	affine hull
$\mathrm{CH}(S)$	convex hull of $S \subseteq \mathbb{R}^d$
$\mathrm{CH}(R\|G)$	convex hull of $R \subseteq \mathbb{R}^d$ relative to $G \subseteq \mathbb{R}^d$
$\mathrm{TH}(R\|G)$	tight hull of $R \subseteq \mathbb{R}^d$ relative to $G \subseteq \mathbb{R}^d$

SUMMARY

Our work applies techniques for blending and tightening solid shapes represented by sets. We require that the output contain one set and exclude a second set, and then we optimize the boundary separating the two sets. Working within that framework, we present mason, tightening, tight hulls, tight blends, and the medial cover, with details for implementation. Mason uses opening and closing techniques from mathematical morphology to smooth small features. By contrast, tightening uses mean curvature flow to minimize the measure of the boundary separating the opening of the interior of the closed input set from the opening of its complement, guaranteeing a mean curvature bound. The tight hull offers a significant generalization of the convex hull subject to volumetric constraints, introducing developable boundary patches connecting the constraints. Tight blends then use opening to replicate some of the behaviors from tightenings by applying tight hulls. The medial cover provides a means for adjusting the topology of a tight hull or tight blend, and it provides an implementation technique for two-dimensional polygonal inputs. Collectively, we offer applications for boundary estimation, three-dimensional solid design, blending, normal field simplification, and polygonal repair. We consequently establish the value of blending and tightening as tools for solid modeling.

PART 1

MASON AND TIGHTENING

CHAPTER 1

INTRODUCTION

Our work details techniques for blending and tightening solid models. Blending, exemplified by the mason operator, replaces sharp or thin features with smooth, thick features. By contrast, tightening minimizes boundary measure. We apply both blending and tightening to sets representing shapes. We often identify one set of fixed-radius balls lying in the interior of the input shape and a second set of fixed-radius balls in its complement. Blending and tightening then optimize the shape boundary lying in the gap between the two sets. As a result, both blending and tightening produce output shapes that resemble the input. In particular, tightening within the gap between the balls eliminates bumps and ripples, simplifying both the surface geometry and its accompanying normal field.

We focus on the use of constrained convexity for blending and tightening. In doing so, we build on the relative convex hull by defining the tight hull, which we design to contain one set and exclude a second while treating the two sets symmetrically. The tight hull consistently reduces boundary measure, simplifies the boundary away from constraints, and introduces topological choices for organizing components of the input. In conjunction with sets of balls lying in the input and its complement, the tight hull blends features and simplifies normals.

Our techniques offer a variety of applications. As an example, we outline a means of boundary measure estimation, which enables us to accurately reconstruct an anatomical structure from medical image data. We also facilitate the design of solid models composed of developable surface patches in the interest of cost-effective manufacturing. In addition, we consider the repair of polygonal models. The fundamental problems that we address provide utility throughout significant aspects of modeling.

In doing so, we draw on the fields of computer graphics, computational geometry, and differential geometry, with significant applications to biomedical imaging.

CHAPTER 2

DOCUMENT STRUCTURE

2.1. Division into parts

This dissertation is divided into two parts:

Part I. After our introduction and outlined document structure, we review our established, peer-reviewed work. The mason operator uses tools from mathematical morphology to make an input shape nearly regular while handling its interior and complement symmetrically. In our work, r-tightening uses mean curvature flow to reduce the input's boundary measure subject to constraints while establishing a mean curvature bound. We then review prior art to establish a context for Part II.

Part II. We detail our recent work, focusing on tight hulls. The tight hull generalizes the convex hull by excluding one set while including another, treating the included and excluded sets symmetrically. Tight blends use tight hulls to recreate aspects of r-tightening and simplify normal fields. After using the medial cover to address topological variations on the tight hull, we discuss implementation and applications, and then we conclude.

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.

Mason. The mason operator uses mathematical morphology to generate a nearly regular output shape. We define r-regularity in terms of balls of radius r in the shape's interior and complement, leading to definitions of the core inside the shape, the anticore in the shape's complement, and the mortar separating the core and anticore. We use compositions of opening and closing operators to generate choices for replacing each component of the mortar, and in each case we choose to minimize the symmetric difference between our input and output. Representing inputs and outputs with either pixels or voxels, we use Euclidean distance mapping and traversal to implement mason with linear time complexity.

Tightening. We define r-tightening, which is a set that locally minimizes boundary measure given that it contains the core and excludes the anticore of an input shape. To minimize boundary measure, we apply mean curvature flow to deform an initial shape. Given that the core and anticore are composed of balls of radius r, the mean curvature of the subset of the boundary supported on the core and anticore ranges from $-\frac{1}{r}$ to $\frac{1}{r}$, while the mean curvature of the subset of unsupported boundary is zero. Consequently, r-tightening guarantees bounded mean curvature. We describe implementations of r-tightening with pixels and voxels using level set flow, with accelerations from narrow band flow and downsampling.

Fairing, blurring, blending. We present aspects of fairing, blurring, and blending as aspects of prior art in solid modeling. Fairing involves surface optimization with energy functionals such as Willmore flow and minimum curvature variation flow; blurring involves modulating eigenfunction amplitudes in mesh signal processing and solid lowpass filtering; and blending involves set theoretic operations in medial axis trimming, relative blending, and the relative convex hull. We characterize these techniques as encompassing a variety of approaches to solid modeling.

2.3. Part II

Tight hulls. 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 hulls 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.

Tight blends. Revising our r-tightening material from Chapter 4 in Part I, we define tight r-blends in terms of tight hulls and morphological opening, and we illustrate how in two dimensions the r-blend of a set has a curvature between $-\frac{1}{r}$ and $\frac{1}{r}$. We describe the deformation of the r-blend of a polygon as we continuously increase the value of r, claiming that the total variation in the normal field of an r-blend monotonically decreases over intervals of r where the blend varies continuously. Combining this result with the curvature-limiting properties of tight blends, we conclude that tight blends simplify normal fields.

Medial cover. 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 that the medial cover, like the tight hull, is symmetric with respect to set complement, while the relative convex hull is not.

Implementation. We explain how to construct two-dimensional tight covers, where the topology of a particular cover depends on our triangulation of the gap between two polygonal sets. We eliminate the variability due to different triangulations by adding 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 produce an approximate tight blend by morphologically 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 tight blends can function as symmetric solid blends, and we suggest that they may prove useful in the multiscale analysis and simplification of normal fields.

Conclusion. We begin by summarizing our contributions in Chapters 3-10. We then discuss the context and implications of our work, emphasizing the tight hull's significance. We conclude by outlining future work.

CHAPTER 3

MASON

3.1. Introduction

On a broad level, our goal involves regularizing complex shapes [**2, 62**]. A closed shape with a manifold boundary is r-regular if and only if both the shape's interior and its complement are equal to unions of open balls of radius r. Consequently, an r-regular shape has guaranteed measures of thickness and smoothness, with no cusps, cracks, or branches. Because regularizing a shape can significantly adjust its geometry, we pursue inducing more subtle changes to geometry to produce nearly regular shapes.

We situate regularity in the context of mathematical morphology, which is well-established. In mathematical morphology, we manipulate sets of balls with a variety of operators. For instance, to grow a shape by radius r we take the union of balls of radius r with centers lying in the shape. To shrink by r we subtract balls with centers lying in a shape's complement. Opening removes material by taking a union of balls in a shape, and closing adds material by taking the complement of a union of balls in its complement [**59, 60**]. We compose opening and closing to produce nearly regular results.

Irregular points constitute the mortar. The mortar is a set of details and features that lie along a shape's boundary, separating the core of points inside the shape from the anticore in its complement. To define the mason operator, we identify components of the mortar and replace each with a component from either opening and then closing or closing and then opening [**74**]. We choose the component that yields the smallest symmetric difference with the input shape, so mason treats points in a shape and its

complement symmetrically. As a result, mason produces nearly regular output while matching the input geometry.

We begin with definitions from mathematical morphology, including our introductions of regularity and the mortar. We next define mason, followed by our two-dimensional and three-dimensional implementations. We discuss possible future work and conclude with our contributions.

3.2. Mathematical morphology

3.2.1. Introduction. We present a framework for the tools from mathematical morphology **[35, 62]** that we use to construct regularity, mortar, and mason. We begin with definitions for balls and sets of balls, then define topological opening and closing. We next use balls to define growing and shrinking as well as morphological opening and closing. We conclude with regularity, core, anticore, and mortar.

3.2.2. Balls and sets of balls. We define a ball B in d-dimensional Euclidean space $\mathbb{R}^d$ using a point $p \in \mathbb{R}^d$ and a radius $r \in \mathbb{R}^>$:

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

The ball $B_r(p)$ consists of all points q with a distance less than r from p, so $B_r(p)$ is open. For an open ball b, let $\rho(b)$ denote the radius of b and let $\pi(b)$ denote its center.

We denote $\mathbb{B}$ as the set of open balls in $\mathbb{R}^d$. Then $\mathbb{B}^S$ is the set of open balls lying in S:

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

Meanwhile, $\mathbb{B}(S)$ is the set of open balls with centers lying in S:

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

Taking $\mathbb{B}_r$ as the set of open balls of radius r,

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

We can combine radius values with symbols for $\mathbb{B}^S$ and $\mathbb{B}(S)$, as in $\mathbb{B}_r^S$ and $\mathbb{B}_r(S)$.

3.2.3. Topological opening and closure. We next define two topological operators. The opening of $S \subseteq \mathbb{R}^d$ is the open set denoted S° as the union of open balls in S:

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

To define closing, we next define complement. The complement of set $X \subseteq \mathbb{R}^d$ is denoted X^c, where $X^c := \mathbb{R}^d \setminus X$ is defined by set difference. Then we obtain the closure of $S \subseteq \mathbb{R}^d$, denoted S^-, by combining two complement operations with opening:

$$S^- := S^{c\circ c}$$

A set's closure contains its boundary, but its opening does not. Denoting the boundary operator with the symbol ∂,

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

As a result, $\partial S \subseteq S^-$ and $\partial S \cap S^\circ = \emptyset$.

3.2.4. Growing and shrinking. The growing of a set S by distance r is denoted by $S \uparrow_r$, which is an open set equal to the union of open balls with radius r whose centers lie in S:

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

Growing S by distance r adds to S the points within distance r of S. Denoting the distance from a point q to S as $d(q, S)$,

$$S \uparrow_r = \left\{ q \in \mathbb{R}^d \mid d(q, S) < r \right\}$$

The convex curves added to S through growing by r have curvature less than or equal to $1/r$. This is because no set of balls centered on the boundary of S can yield a convex curvature greater than $1/r$.

The shrinking of a set S by distance r is a closed set denoted $S \downarrow_r$. The relationship between growing and shrinking is comparable to the relationship between opening and closing. In particular, shrinking can be expressed as combinations of growing and complement:

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

Shrinking S by distance r yields the points at distance greater than or equal to the complement of S:

$$S \downarrow_r = \left\{ q \in \mathbb{R}^d \mid d(q, S^c) \geq r \right\}$$

Comparable to growing, shrinking by r introduces concave curves with curvature greater than $-1/r$. The balls centered on the boundary of S that are removed cannot leave a curvature less than $-1/r$.

3.2.5. Morphological opening and closing. Opening a set S by radius r, denoted $S\circ_r$, takes the union of open balls of radius r contained in S:

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

As in growing, the convex curves introduced by opening with r have curvature less than or equal to $1/r$. As in the relationship between topological opening and closing and the relationship between growing and shrinking, the morphological closing of a set S by radius r, denoted $S\bullet_r$, is equal to combinations of opening and complement:

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

The concave curves introduced by closing have curvature greater than or equal to $-1/r$ [**15**].

3.2.6. Regularity. The regularity of a point p with respect to $S \subseteq \mathbb{R}^d$, denoted $\mathcal{R}^S(p)$, is the radius of the largest open ball in S whose closure contains p. If p is not contained by the closure of an open ball in S, the regularity of p with respect to S is equal to zero:

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

If a point $p \in S$ has a regularity greater than or equal to r with respect to S, we refer to p as r-inner-regular with respect to S. Similarly, if a point $p \in S^c$ has a regularity greater than or equal to r with respect to S^c, we refer to p as r-outer-regular with respect to S. We refer to the set S as r-inner-regular if every point p in S is r-inner-regular, so $S^\circ = S\circ_r$. We likewise refer to S as r-outer-regular if every point p in S^c is r-outer-regular, so $S^{c\circ} = S^c\circ_r$. A set S is r-regular if S is both r-inner-regular and r-outer-regular.

3.2.7. Core, anticore, and mortar. The set $S\circ_r$ constitutes the r-core of S, while $S^c\circ_r$ constitutes the r-anticore of S. Points that lie neither in the r-core nor the r-anticore constitute the r-mortar. We denote the mortar for radius r and set S as $\mathcal{M}_r S$, where $\mathcal{M}_r S$ is the intersection of the r-closing of S and the r-closing of the complement of S:

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

The concave curves along the boundary of the r-mortar have curvature between $-1/r$ and 0. We also refer to points in the interior of the r-mortar as r-irregular.

3.3. Mason

3.3.1. Combinations of opening and closing. In this section, we define the mason operator. Just as the morphological opening of a set S is inner-regular and the morphological closing of S is outer-regular, the combination $S \bullet_r \circ_r$ of closing and opening is r-inner-regular, while the combination $S \circ_r \bullet_r$ of opening and closing is r-outer-regular. In practice, both $S \bullet_r \circ_r$ and $S \circ_r \bullet_r$ are almost everywhere r-regular because they eliminate most irregular points. As a result, we exploit $S \circ_r \bullet_r$ and $S \bullet_r \circ_r$ to maximize mason's regularity.

Note that irregularities in two dimensions left by combinations of opening and closing are either cusps or bridges. The cusps and bridges in $S \circ_r \bullet_r$ lie in $S \circ_r \bullet_r$, while the cusps and bridges in $S \bullet_r \circ_r$ lie in the complement $S \bullet_r \circ_r{}^c$.

3.3.2. Asymmetry. We defined the r-mortar of set S as the intersection of $S\bullet_r$ and $S^c\bullet_r$ because it highlights a symmetry between a set and its complement. The r-mortar is also equal to the difference between r-closing and r-opening:

$$\mathcal{M}_r S = S \bullet_r - S\circ_r$$

Opening a set empties the mortar, while closing the set fills it. Closing after opening simply refines the boundary of the opened set. Similarly, opening after closing refines closing. As a result, closing and then opening is biased to increase a set's area or volume, while opening and then closing is biased to decrease area or volume.

The operators $S\circ_r$, $S\bullet_r$, $S\circ_r\bullet_r$, and $S\bullet_r\circ_r$ are asymmetric with respect to the set S and its complement S^c. For instance, opening a two-dimensional set S decreases its area, while opening the complement of S and then taking the complement increases area. By contrast, a symmetric operator yields the same result when applied to S as when it is applied to S^c and followed by its complement. Symmetry is desirable for a variety of reasons. For instance, it handles cracks and branches as similar features. We consequently incorporate symmetry into our design of mason.

3.3.3. Local decisions. To obtain mason, we extract regions from the r-mortar of an input set S, and we choose each region for mason from either $S\circ_r\bullet_r$ or $S\bullet_r\circ_r$ to minimize its difference in area or volume from S. We take the regions C_rS as the set of closed connected components from the interior of the mortar. We define C_rS in terms of κ, which denotes the set of connected components from an input set:

$$C_rS := \left\{\rho^- \mid \rho \in \kappa\left((\mathcal{M}_rS)^\circ\right)\right\}$$

Letting T_rS denote the mason of S by radius r, the subset of T_rS outside the union of ∂S with union C_rS is equal to the r-core of S. We then identify a particular region in C_rS with an index i as C_r^iS. We take the intersection of C_r^iS with T_rS equal to either $C_r^i\left(S\circ_r\bullet_r\right)$ or $C_r^i\left(S\bullet_r\circ_r\right)$, depending on which of $C_r^i\left(S\circ_r\bullet_r\right)$ or $C_r^i\left(S\bullet_r\circ_r\right)$ has a minimal Hausdorff measure H^d of set difference $\setminus$ with C_r^iS in $\mathbb{R}^d$:

$$C_r^iS\cap T_rS = \begin{cases} C_r^i\left(S\circ_r\bullet_r\right) & \mathrm{H}^d\left(C_r^iS\setminus C_r^i\left(S\circ_r\bullet_r\right)\right) \leq \mathrm{H}^d\left(C_r^iS\setminus C_r^i\left(S\bullet_r\circ_r\right)\right) \\ C_r^i\left(S\bullet_r\circ_r\right) & \mathrm{H}^d\left(C_r^iS\setminus C_r^i\left(S\bullet_r\circ_r\right)\right) < \mathrm{H}^d\left(C_r^iS\setminus C_r^i\left(S\circ_r\bullet_r\right)\right) \end{cases}$$

By choosing components of $C_r^i\,(S \circ_r \bullet_r)$ and $C_r^i\,(S \bullet_r \circ_r)$ to minimize the set difference between mason and $C_r^i\,(S)$, mason is essentially symmetric because it makes essentially the same choices to generate T_rS and $(T_rS^c)^c$. By combining combinations of opening and closing in components of the mortar, mason also eliminates most irregularities. Consequently, mason offers attractive properties.

3.4. Implementation

3.4.1. Euclidean distance mapping. To implement mason, we use Euclidean distance mapping [**22**] to compute morphological operations for binary images represented as pixels or voxels. The algorithms for pixels and voxels are largely identical, so we refer to pixels. Each pixel stores a two-element vector that points from a pixel to the closest pixel of opposite color, yielding a vector from a white pixel to a black pixel or a vector from a black pixel to a white pixel. The distance mapping algorithm propagates vector updates by sweeping from left to right and proceeding, consecutively, from row to row. At a pixel, the update uses an eight-point set of neighbors. The sweep then updates from right to left and bottom to top. Together, the two sweeps yield distance errors bounded in practice by 0.076 pixels.

3.4.2. Morphology from Euclidean distance mapping. In our discussion of our implementation, we use S to denote a set of black pixels or voxels in an image. Similarly, $S \uparrow_r$ denotes the result of adding pixels to S that lie within r units times the distance between two pixels. We likewise use $S \downarrow_r$, $S\circ_r$, $S\bullet_r$, $S \circ_r \bullet_r$, $S \bullet_r \circ_r$, $\mathcal{M}_rS$, and $\mathcal{T}_rS$ for denoting operators applied to images.

To compute $S \uparrow_r$, we take a threshold of the Euclidean distance map of S to make black the white pixels near the input set's boundary. Similarly, we shrink the input set with $S \downarrow_r$ by using a different threshold to make black pixels near its boundary white. By combining growing and shrinking, we obtain $S\circ_r$ and $S\bullet_r$, which we combine to

obtain $S \circ_r \bullet_r$ and $S \bullet_r \circ_r$. The mortar is then the difference in pixels between closing and opening.

3.4.3. Identifying mortar regions. We identify regions in the mortar by marking core and anticore pixels with a constant label while mortar is unlabeled. We traverse the image until we identify an unlabeled mortar pixel and then we assign it a new label. We next begin a breadth-first traversal, applying the same label to the pixel's unlabeled neighbors in the mortar and adding them to a queue. Once we complete iterating over the queue, all of the pixels in the same region share the same identifying label.

3.4.4. Selecting regions for mason. For each region in the mortar of an image, we identify the pixels from $S \circ_r \bullet_r$ and the pixels from $S \bullet_r \circ_r$. We choose for mason the pixels that yield the least change in the number of pixels in the region from $\mathcal{M}_r S$.

3.4.5. Complexity. Euclidean distance mapping for an image with n pixels has a time complexity in $O(n)$ because it simply traverses pixels twice. Similarly, the morphological operations $S \uparrow_r$, $S \downarrow_r$, $S\circ_r$, $S\bullet_r$, $S \circ_r \bullet_r$, $S \bullet_r \circ_r$, and $\mathcal{M}_r S$ have $O(n)$ time complexity because they require no more than traversing pixels after distance mapping. Identifying mortar regions in an image with n pixels also has $O(n)$ time complexity, because each pixel in a mortar region is only enqueued and dequeued once, while it is accessed by no more than its number of neighbors. Selecting regions for mason also only requires linear time to compare the pixel values of S, $S \circ_r \bullet_r$, and $S \bullet_r \circ_r$ in the regions from $C_r S$, yielding an overall $O(n)$ complexity.

3.4.6. Two-dimensional results. We produce three images, each yielding $S \circ_r \bullet_r$, $S \bullet_r \circ_r$, and $\mathcal{T}_r S$ from input image S. For each image, we compute the area of the image and the numbers of pixels changed, black components, white components, inner-irregular pixels, and outer-irregular pixels. Computing our two-dimensional images requires negligible time.

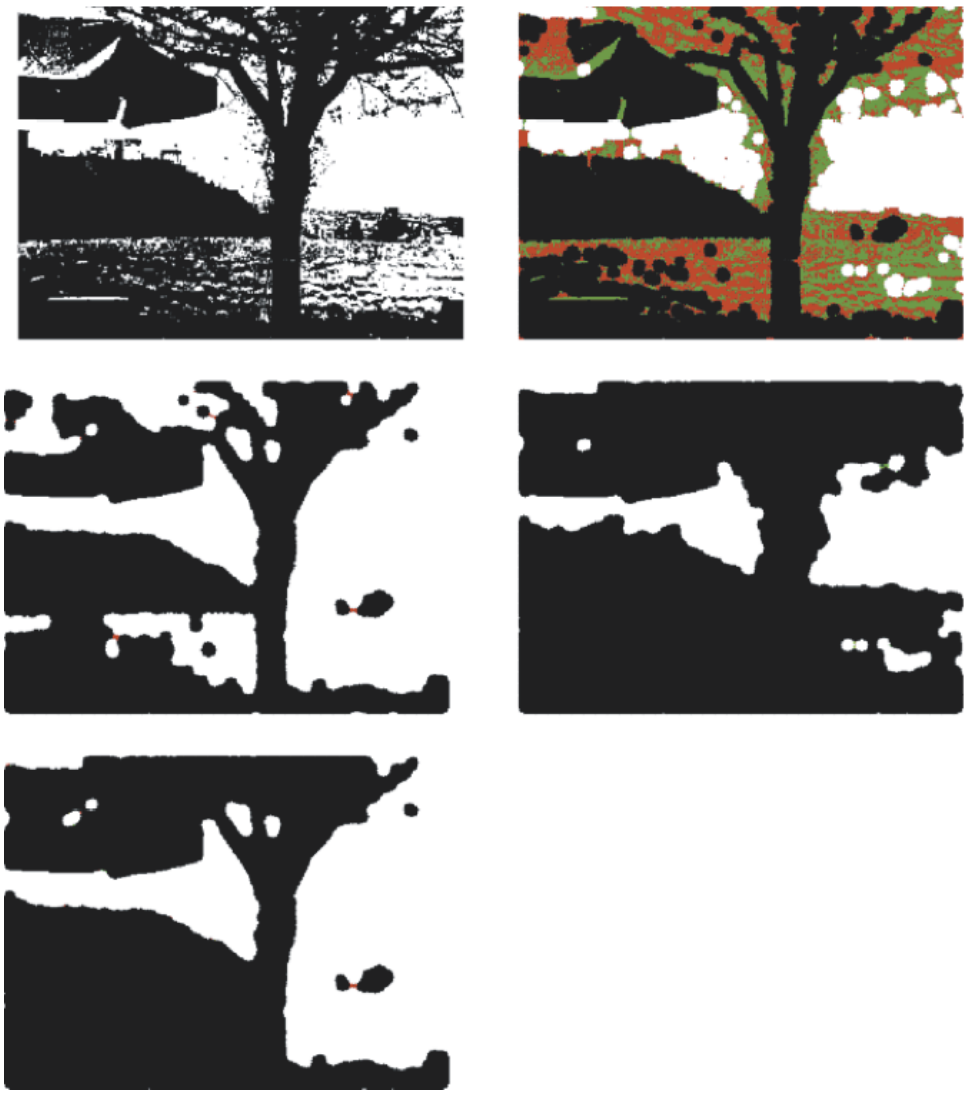


FIGURE 3.1. (Top left) The input set S shown with black pixels. (Top right) The inner-irregular points from S are red, the outer-irregular points are green, and the union of the red and the green points are the mortar. The core is black and the anticore is white. (Middle left) The regular points of $S \circ_5 \bullet_5$ are black, the regular points of the complement $S \circ_5 \bullet_5{}^c$ are white, and the inner-irregular points are red. (Middle right) The regular points of $S \bullet_5 \circ_5$ are black, and the outer-irregular points are green. (Lower left) Inner-irregular points of mason are red and outer-irregular points in green.

Our first image exhibits mason's advantages (Figure 3.1, Table 1.) Illustrating a landscape, the image is 400 by 300 pixels with radius $r = 5$. The output results $S \circ_r \bullet_r$, $S \bullet_r \circ_r$, and $\mathcal{T}_r S$ all reduce the numbers of black and white components as well as the numbers of inner-irregular and outer-irregular pixels. Mason output is

TABLE 1. Columns list the area of each image, the number of pixels changed relative to S, the number of black components, the number of white components, the number of inner-irregular points, and the number of outer-irregular points.

	Area	Change	Black	White	Inner	Outer
S	78,195	N/A	527	711	18,676	18,641
$S \circ_5 \bullet_5$	62,808	19,301	6	23	107	0
$S \bullet_5 \circ_5$	95,472	18,573	1	33	0	43
$\mathcal{T}_5 S$	75,674	15,071	3	27	48	13

TABLE 2. Quantitative results for Figure 3.2.

	Area	Change	Black	White	Inner	Outer
S	36,967	N/A	98	177	7,970	7,312
$S \circ_5 \bullet_5$	31,873	9,218	8	17	203	0
$S \bullet_5 \circ_5$	41,842	8,523	5	20	0	89
$\mathcal{T}_5 S$	38,013	7,922	5	22	203	108

neither inner-regular nor outer-regular, but it is nearly regular, in contrast to the input image. Closing and then opening fills components of $\mathcal{M}_r S$, while opening and then closing clears them out. Mason chooses the component from $S \circ_r \bullet_r$ to clear the predominantly white pixels of water to the right of the tree, but in contrast mason chooses $S \bullet_r \circ_r$ to fill the predominantly black region at the upper left of the tree. Mason's pattern of choices results in less area change than $S \circ_r \bullet_r$ or $S \bullet_r \circ_r$, highlighting mason's symmetry.

Our second image presents an example where $S \circ_r \bullet_r$, $S \bullet_r \circ_r$, and $\mathcal{T}_r S$ yield similar performance (Figure 3.2, Table 2.) The data in the image is from a slide of lung cells **[26]**, consisting of 300 by 225 pixels with $r = 5$. The connected components of the mortar contain comparable quantities of inner-irregular and outer-irregular pixels, so the patterns of filling and clearing exhibited by $S \circ_r \bullet_r$, $S \bullet_r \circ_r$, and $\mathcal{T}_r S$ yield comparable changes in area. All three techniques similarly reduce the numbers of components and the numbers of irregular pixels, so mason does not exhibit a strong advantage.

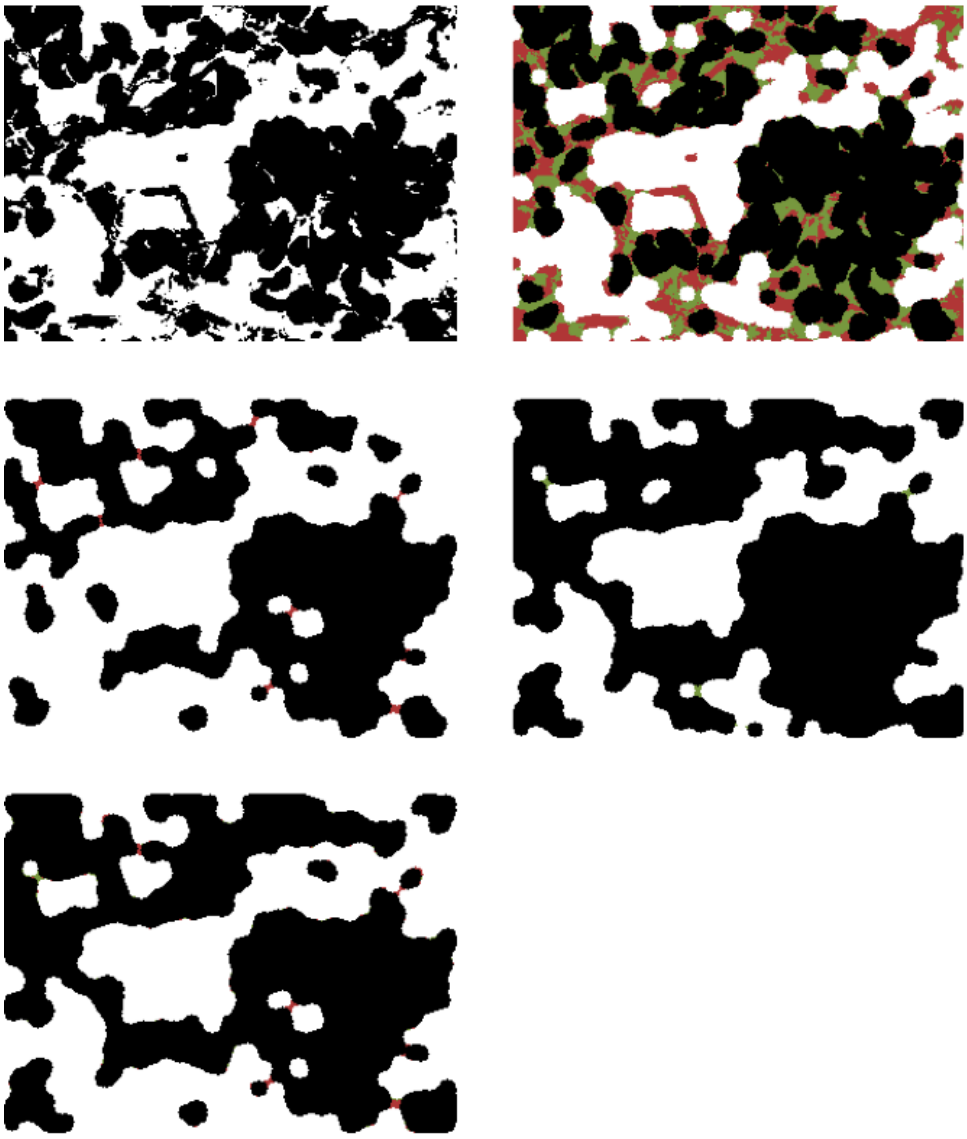


FIGURE 3.2. (Top left) Input set S. (Top right) $\mathcal{M}_5 S$ (Middle left) $S \circ_5 \bullet_5$ (Middle right) $S \bullet_5 \circ_5$ (Lower left) $\mathcal{T}_5 S$

TABLE 3. Quantitative results for Figure 3.3.

	Area	Change	Black	White	Inner	Outer
S	128,956	N/A	2,212	3,218	101,717	79,769
$S \circ_5 \bullet_5$	29,993	100,969	64	6	679	0
$S \bullet_5 \circ_5$	204,213	79,225	9	145	0	154
$\mathcal{T}_5 S$	203,329	78,923	3	145	15	107

Our final image demonstrates that under some circumstances $S \circ_r \bullet_r$, $S \bullet_r \circ_r$, and $\mathcal{T}_r S$ all generate limited results (Figure 3.3, Table 3.) The image of a galaxy with 620 by 430 pixels **[66]** also has radius $r = 5$. The mortar contains a single large component

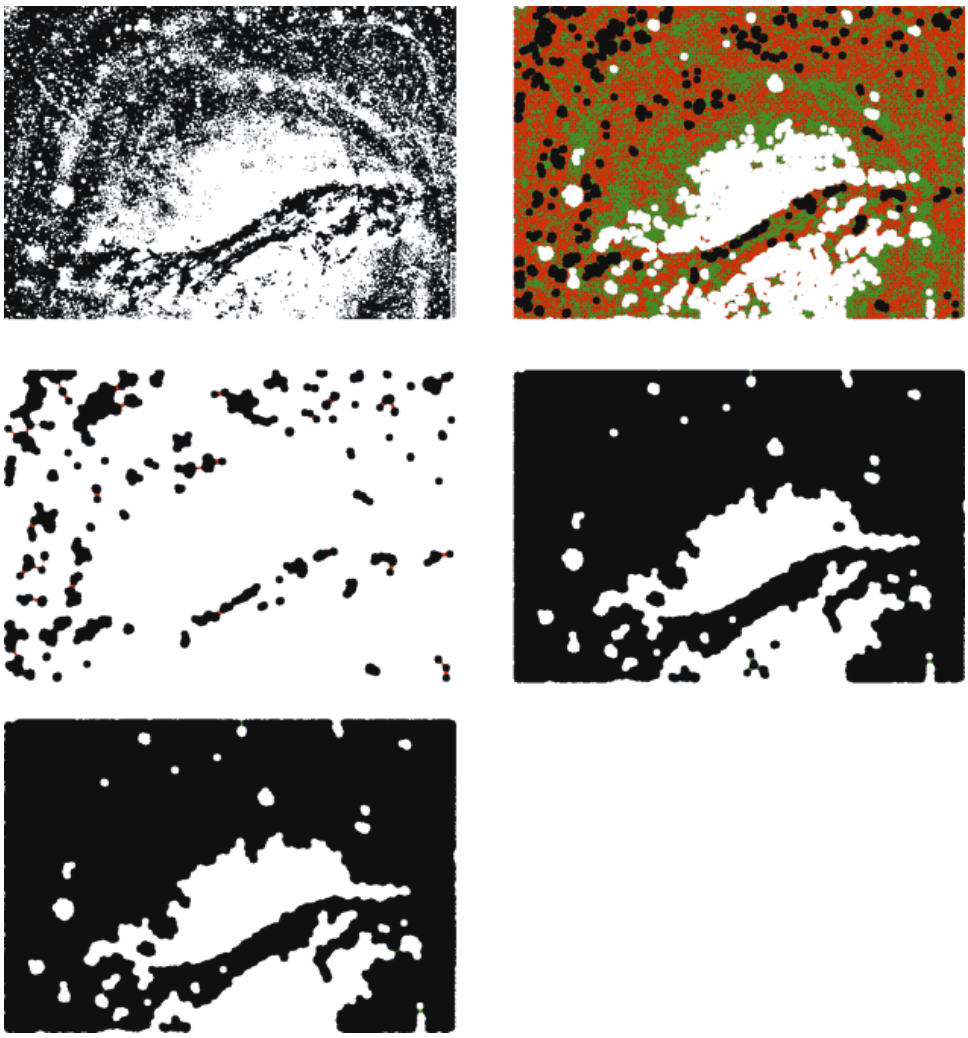


FIGURE 3.3. (Top left) Input set S. (Top right) $\mathcal{M}_5 S$ (Middle left) $S \circ_5 \bullet_5$ (Middle right) $S \bullet_5 \circ_5$ (Lower left) $\mathcal{T}_5 S$

that consists of a combination of some areas that are predominantly black while other areas are predominantly white. Mason makes a single choice for the large component, so $\mathcal{T}_r S$ matches $S \bullet_r \circ_r$ by filling the component while $S \circ_r \bullet_r$ clears it. Because of the combination of black and white pixels in the component, $S \bullet_r \circ_r$, $S \circ_r \bullet_r$, and $\mathcal{T}_r S$ all result in a large area change. Mason can only exhibit advantage when provided significant choice.

3.4.7. Three-dimensional results. We show the results of processing four different solids represented as binary volumes. The first binary volume spans 425 x 150 x 150 voxels (Figure 3.4.) On the right of the upper left volume, 27 overlapping balls

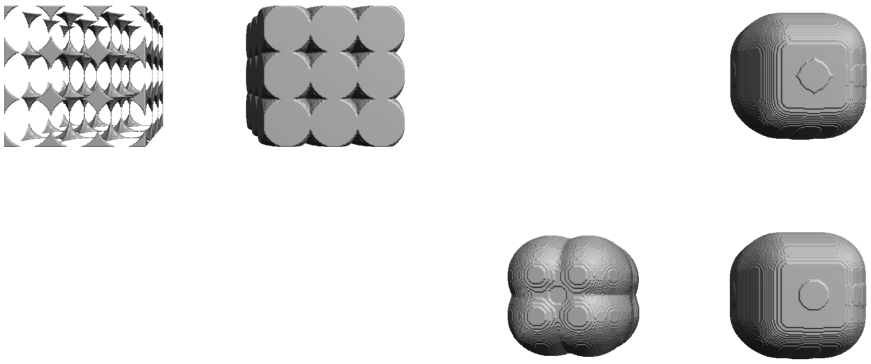


FIGURE 3.4. (Upper left) Input set S. (Upper right) $\mathcal{T}_{50}S$ (Lower left) $S \circ_{50} \bullet_{50}$ (Lower right) $S \bullet_{50} \circ_{50}$

fill most of a cube, while the difference between the cube and the balls appears on the left. The solid contains 3,375,000 voxels, and we apply a radius $r = 50$. First opening and then closing clears both the left and right sides, while closing and then opening fills them. Mason, by contrast, fills the balls on the right and clears the left. The symmetric differences that we obtain are 3,375,000 voxels for $S \circ_r \bullet_r$, 3,068,335 voxels for $S \bullet_r \circ_r$, and $95,308$ for $\mathcal{T}_r S$. As is visually apparent, mason more closely represents the input set S than does $S \circ_r \bullet_r$ or $S \bullet_r \circ_r$.

The second volume spans 363 x 231 x 514 voxels, and we process using radius $r = 32$ (Figure 3.5.) We extract isosurfaces from signed distance fields to illustrate that $S \circ_r$ is a union of spheres, while the complement $S \bullet_r{}^c$ is a union of spheres. Mason chooses essentially the same mortar components as $S \bullet_r \circ_r$, resulting in a symmetric difference of $1,870,597$ voxels for $\mathcal{T}_r S$, nearly equal to the symmetric difference of 1,872,535 voxels for $S \bullet_r \circ_r$. The symmetric difference of $S \circ_r \bullet_r$ is significantly larger than both, amounting to 2,336,582 voxels. While $S \bullet_r \circ_r$ may match mason's performance, neither $S \bullet_r \circ_r$ no $S \circ_r \bullet_r$ can achieve a smaller symmetric difference.

The third volume highlights how $S \circ_r \bullet_r$ and $S \bullet_r \circ_r$ can be nearly identical, so mason yields a nominal difference (Figure 3.6.) The fourth volume shows the effects of

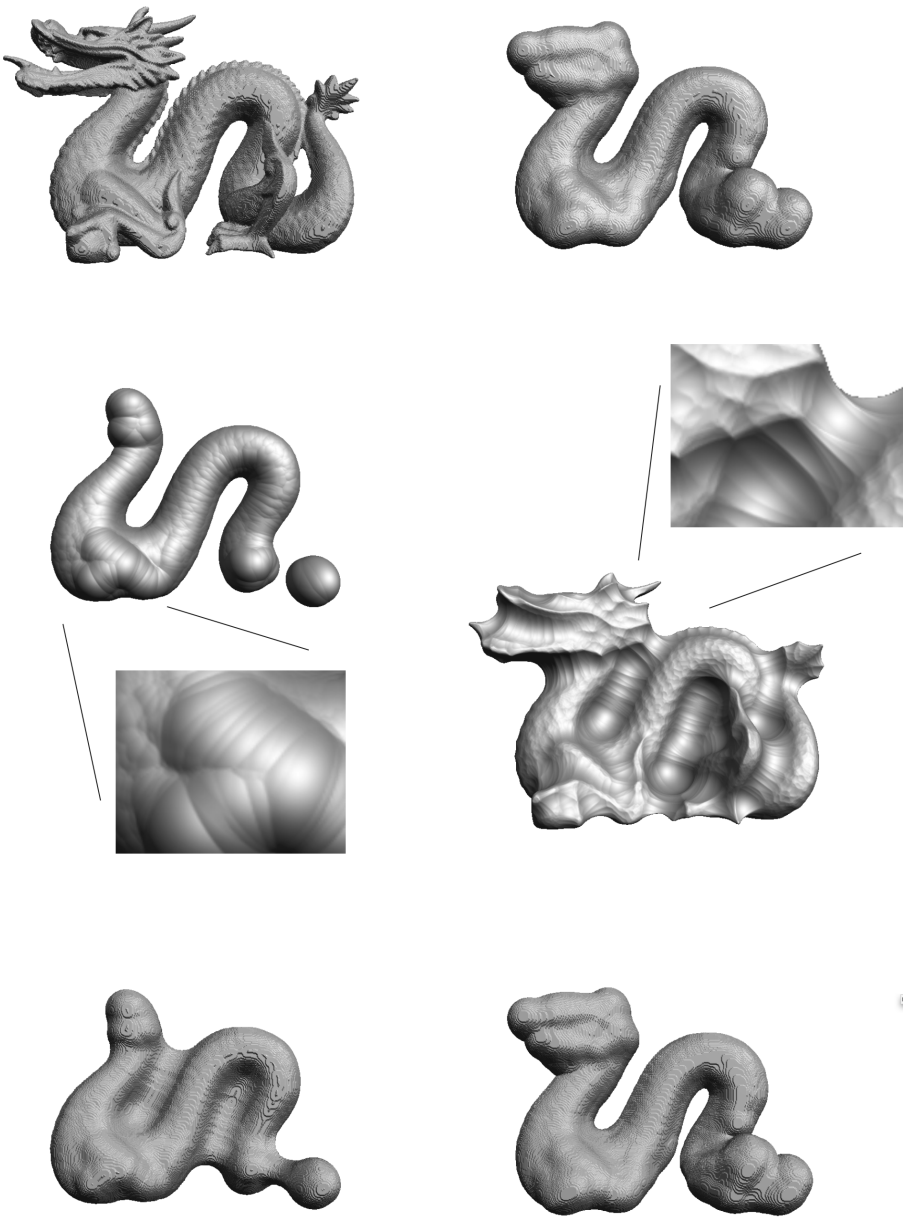


FIGURE 3.5. (Upper left) Input set S. (Upper right) $\mathcal{T}_{32}S$ (Middle left) An inset illustrating that $S\circ_{32}$ is a union of spheres. (Middle right) An inset illustrating that $S \bullet_{32}{}^{c}$ is a union of spheres. (Bottom left) $S \circ_{32} \bullet_{32}$ (Bottom right) $S \bullet_{32} \circ_{32}$

varying the radius used in mason, demonstrated by 16, 32, 48, and 64 voxels (Figure 3.7.)

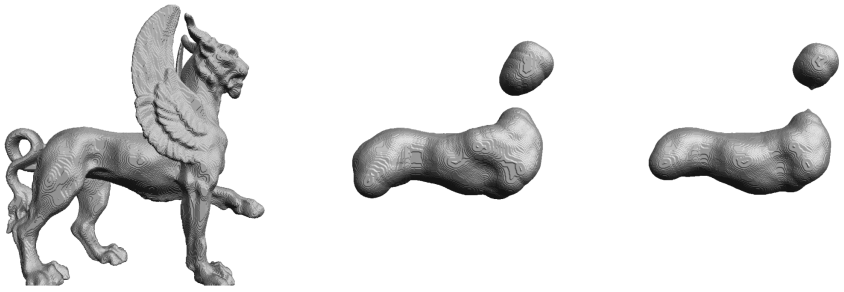


FIGURE 3.6. (Left) Input set S. (Middle) $S \bullet_r \circ_r$ (Right) $S \circ_r \bullet_r$

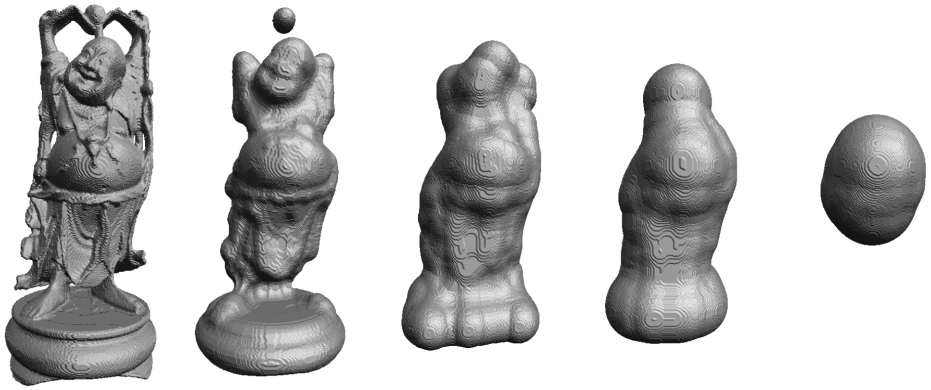


FIGURE 3.7. To the left, we show the input set S. Progressively to the right, we show $\mathcal{T}_{16}S$, $\mathcal{T}_{32}S$, $\mathcal{T}_{48}S$, and $\mathcal{T}_{64}S$.

The time required by computing our results with binary volumes is dominated by the time required by distance transforms. The time complexity is linear in the number of voxels, requiring roughly 15 seconds to process 1,000,000 voxels using an 800 MHz Powerbook G4. The input volumes range from 10 to 56 million voxels, but the growing operation requires padding each volume by a corresponding radius. The padded volumes range from 33,000,000 to 94,000,000 voxels, each requiring 8-24 minutes for a distance transform. Processing $S\circ_r$ and $S\bullet_r$ require two transforms, $S \circ_r \bullet_r$ and $S \bullet_r \circ_r$ require four transforms, and $\mathcal{T}_r S$ requires eight transforms. As a result, computing the transforms for mason requires 1-3 hours, and overhead for selecting mortar regions requires 5-28 minutes.

3.5. Discussion

In practice, the mason operator yields a nearly regular output, while its regularity is not guaranteed. Because we consider regularity desirable, our work on mason raises the issue of how to design a regularizing operator. In particular, we ask how we can minimize change in the regularized output set's geometry, given that in some circumstances regularity cannot be obtained by exclusively modifying the mortar.

At a point p on the boundary of an r-regular set, the curvatures through p range from $-\frac{1}{r}$ to $\frac{1}{r}$. Because mason is nearly regular, the curvatures are limited by $\left|\frac{1}{r}\right|$ at most points. Consequently, we ask whether we can generate a set with a curvature-limited boundary that lies in the mortar. In two dimensions, this entails eliminating cusps, so we ask if we can eliminate cusps using limited time complexity.

In both two and three dimensions, we directly implement set theoretic operations for input data by comparing pixel by pixel or voxel by voxel. We therefore ask if we can compute mason using boundary representations as input and output. Boundary representations promise increased numerical precision, but the algorithms used to manipulate them are more complex.

3.6. Contributions

As our contributions in this chapter, we define regularity, mortar, and the mason operator, as well as mason's two- and three-dimensional implementations. We define the r-regularity of a point with respect to a set as the radius of the largest ball lying in the set and containing the point. As an outcome, we define r-inner-regular, r-outer-regular, r-inner-irregular, and r-outer-irregular points in terms of their regularity, and we define sets as r-inner-regular, r-outer-regular, or r-regular depending on the regularity values of points with respect to sets and their complements. Qualitatively, irregularities identify cracks, cusps, and other fine details, while regular shapes are smooth and thick, with bounded curvature. In contrast to the distance transform,

regularity indicates the distribution of space in relation to a set. As a result, regularity offers a promising characterization of shape.

Mortar functions as a tolerance zone consisting of irregular points, separating the inner-regular core from the outer-regular anticore. Because we express the r-mortar of S as the intersection of the closing $S\bullet_r$ and the closing of the complement $S^c\bullet_r$, mortar is analogous to the boundary operator, which intersects the closure S^- with the closure of the complement S^{c-}. Consequently, mortar preserves the regular parts of a shape and its complement while permitting changes to the details lying along the shape's boundary. Mortar therefore behaves differently from a tolerance zone such as a bound on Hausdorff error. A bound on Hausdorff error allows changes along a shape's boundary independent of the size of details. A resulting shape may therefore induce more significant changes in appearance and physical properties than variations we obtain in the mortar.

Mason modifies the contents of the mortar by replacing each component of $\mathcal{M}_r S$ with a component from either $S \circ_r \bullet_r$ or $S \bullet_r \circ_r$. By making a sequence of choices, mason produces nearly regular output, and its symmetric difference with S is less than or equal to the symmetric differences of $S \circ_r \bullet_r$ and $S \bullet_r \circ_r$ with S. We also demonstrate that mason is effectively symmetric, treating points in S and S^c as possessing identical properties, so $\mathcal{T}_s S = (\mathcal{T}_r S^c)^c$. Mason typically yields a greater visual match with S than either $S \circ_r \bullet_r$ or $S \bullet_r \circ_r$ while retaining desirable qualities shared by both.

We demonstrate viable two- and three-dimensional implementations of mason, where in two dimensions we represent input using binary images composed of pixels, while in three dimensions we use voxels. Our morphological operators involve Euclidean distance transforms, requiring linear time. The additional operations for mason also require no more than linear time, so mason is applicable to large inputs in practice.

Together, regularity, mortar, and mason provide tools for identifying and blending features from solids.

CHAPTER 4

TIGHTENING

4.1. Introduction

We define the r-tightening of a set $S \subseteq \mathbb{R}^d$, where the boundary of the r-tightening separates the union of balls of radius r contained in S from the union of balls of radius r contained in the complement of S [**75**]. The r-tightening's boundary locally minimizes a specific measure, and as a consequence the mean curvature at every boundary point ranges from $-\frac{1}{r}$ to $\frac{1}{r}$. This bound on the curvature of an r-tightening is a valuable property that is not guaranteed by morphological operations such as opening, closing, and mason.

We construct r-tightening by applying mean curvature flow. Using the definitions of opening, closing, mason, mortar, core, and anticore in Chapter 3, we identify an initial set T containing the core and excluding the anticore of the input set S. We displace each boundary point lying in the mortar between the core and the anticore by the normal vector scaled by the mean curvature at that point. The resulting deformation displaces T's boundary along its gradient.

Using two-dimensional images and a combination of three-dimensional triangle meshes and grids, we define level set constrained mean curvature flow. We implement level set flow and two acceleration techniques, yielding practical algorithms.

4.2. Definition and properties

4.2.1. Definition of r-tightening. An r-tightening of a set contains the set's core while excluding its anticore. The boundary components of an r-tightening also have locally minimum measure, which entails minimal length in two dimensions and

minimal area in three dimensions. A boundary's measure is locally minimal in the sense that no deformation of the r-tightening reduces its boundary measure while it remains an r-tightening.

To formalize r-tightening, we define candidate r-tightening in order to establish an r-tightening's minimal boundary. In our definition of r-tightening, we use H^{d-1} to denote the $(d-1)$-dimensional Hausdorff measure [**33**] of the boundary of a d-dimensional set:

Definition. A set $C \subseteq \mathbb{R}^d$ is a *candidate r-tightening* of $S \subseteq \mathbb{R}^d$ if and only if $S \circ_r \subseteq C$ and $S \bullet_r{}^c \cap C = \emptyset$.

A candidate r-tightening set $T \subseteq \mathbb{R}^d$ is an *r-tightening* of $S \subseteq \mathbb{R}^d$ if and only if there is no candidate r-tightening set $C \subseteq \mathbb{R}^d$ such that $\mathrm{H}^{d-1}(C) < \mathrm{H}^{d-1}(T)$ given that C is isotopic to T throughout an isotopy in which C continuously remains a candidate r-tightening.

4.2.2. Mean curvature flow. We use constrained mean curvature flow to construct the boundary of an r-tightening. To deform a surface, mean curvature flow assigns a mean curvature vector to each surface point equal to a normal vector with its length equal to the mean curvature at that point. For a smooth surface, mean curvature at a point is the mean of the principal curvatures at the same point. In three dimensions, for instance, a smooth surface consists of two dimensions with two principal curvatures, so we denote mean curvature with two dimensions by H^2 and principal curvatures by κ_1 and κ_2:

$$H^2 = \frac{1}{2}(\kappa_1 + \kappa_2)$$

In two dimensions, by contrast, a smooth boundary consists of a single dimension with a curvature at each point. The one-dimensional mean curvature H^1 is then equal to the curvature κ.

Mean curvature flow moves each point on a smooth surface along its mean curvature vector. The motion deforms the surface along the gradient of its measure, reducing the measure. As proven by [**32**], an unconstrained boundary curve collapses to a point, while work such as [**25**] extends results to higher dimensions. Given constraints, a surface deforms until it is stable [**21**]. Once it is stable, the surface has minimal measure. Consequently, we obtain r-tightening from constrained mean curvature flow.

4.2.3. Curvature. Under stable, constrained mean curvature flow, mean curvature normals that point into the constraints have lengths determined by the geometry of the constraints supporting them. Unsupported mean curvature normals instead have zero length. In two dimensions, the unsupported boundary consists of line segments, while in three dimensions the unsupported boundary consists of saddles. In either case, the unsupported boundary has zero mean curvature.

An r-tightening's supported boundary is constrained by the complement of the mortar, so that it rests on the mortar's concave boundary. The mean curvature of the mortar's concave boundary ranges from $-\frac{1}{r}$ to 0, but the subset of an r-tightening's boundary supported by the boundary of the core is convex, while the subset supported by the boundary of the anticore is concave. The mean curvature of an r-tightening therefore ranges from $-\frac{1}{r}$ to $\frac{1}{r}$.

The principal curvatures of the supported boundary also range from $-\frac{1}{r}$ to $\frac{1}{r}$. Although the unsupported boundary has zero mean curvature, however, we have no range of values for principal curvatures at unsupported points. In three dimensions, the two principal curvatures have equal magnitudes and opposite signs. Their sum is zero, but their magnitudes may be large.

4.2.4. Topology. Every boundary component of an r-tightening contains at least one component of the core, assuming the boundary component is bounded. This is because a bounded boundary component that does not contain any subset of the core collapses like an unconstrained surface.

There may be more than one r-tightening of the same given set. Our definition requires that no r-tightening can deform into an r-tightening with less boundary measure, provided that we continuously contain the core and exclude the anticore. It may be impossible to deform two r-tightenings into each other because either they are not isotopic or do not continue to be r-tightenings during the isotopy.

If we are given a set S, we can use constrained mean curvature flow to deform a set T that contains the core of S and excludes its anticore. The deformation yields an r-tightening, but which r-tightening it produces depends on the choice of T. In principle, we can choose T to minimize the boundary measure of the resulting r-tightening, but the choice is not generally obvious. We alternatively select mason, so T resembles S. As a result, T yields an r-tightening that also resembles S.

4.3. Two-dimensional implementation

4.3.1. Level set curvature flow. We implement r-tightening in two dimensions by applying level set curvature flow **[63]** to binary images represented by pixels. To deform a binary set, we define a function value at each pixel. Negative values represent the set's interior, while positive values represent its exterior. To displace a boundary interval outward, we decrease values lying along it, and to displace a boundary interval inward, we increase values along it.

We denote the function as Φ, which we initialize to be the signed distance to the boundary of the input binary set. For specific values, $\Phi(t, x, y)$ represents a pixel at coordinates x and y at time t. Function Φ_t denotes the change in Φ at time t. For mean curvature flow, we express Φ_t as the negative product of the mean curvature κ and the magnitude of the gradient of Φ:

$$\Phi_t = -\kappa \left| \nabla \Phi \right|$$

The value of κ is the divergence of the normalized gradient:

$$\kappa = \nabla \cdot (\nabla\Phi / |\nabla\Phi|)$$

The gradient is a normal vector pointing away from the set's interior and toward its exterior. The normalized gradient is then a unit normal vector, so κ is the divergence of local unit normal vectors. Because we compute Φ at grid points, we compute gradients and divergence using finite differences.

To update Φ, we adjust Φ_t using a value $\triangle t$ inversely proportional to the maximum value of $\Phi_t(t, x, y)$ so that the zero level boundary in Φ changes by at most one pixel each iteration:

$$\Phi(t + \triangle t, x, y) = \Phi(t, x, y) + \triangle t \cdot \Phi_t(t, x, y)$$

4.3.2. Constrained mean curvature flow. To compute constrained mean curvature flow, we make no changes to the value of Φ in the core or anticore. As in Subsection 3.2.4, we define the core of binary input set S using the growing and shrinking operations $\uparrow$ and $\downarrow$. The core is then $S \downarrow\uparrow$ and the anticore is $S \uparrow\downarrow^{c}$. We compute $\uparrow$ and $\downarrow$ by thresholding a Euclidean distance transform for each operation. Each operation has linear complexity, so the total of four operations required by the core and anticore also has linear complexity. Once we compute the core and anticore, values of Φ within them are set by the signed distance to S.

4.3.3. Narrow band. We apply two techniques to accelerate level set flow. As our first technique, we exclusively update boundary pixels lying in a narrow band. We begin by building a list of pixels adjacent to pixels of opposite sign, which defines the boundary. We then use a breadth-first search to add all the pixels lying within three units of the boundary. We continue updates within the band as long as the boundary remains within the band. When updates displace the boundary so that it lies outside



FIGURE 4.1. (Left) The input binary image consists of 1102 x 832 pixels. (Right) The tightening of the input with a radius of 40 pixels.

the band, we recompute the band. Using the band improves the performance speed by roughly an order of magnitude.

4.3.4. Downsampled image representation. As our second technique, we downsample the image representation and then iteratively perform level set flow, progressively upsampling the image each time the flow converges. We first downsample the image until the value r for tightening reduces to less than two pixels. Once the flow converges, we upsample by a factor of two and begin the flow again. The number of iterations needed for convergence at the coarsest level depends on the thickness of the mortar, while the number of iterations that perform smoothing remains stable across levels.

4.3.5. Results. Figure 4.1 shows an r-tightening given an image consisting of 1102 x 832 pixels computed with radius $r = 40$. The morphological operations required two seconds, while the downsampled narrow band curvature flow required five seconds for one hundred levels at each of five levels of resolution while using a dual 2 GHz Power PC G5.

In Figure 4.2 , we show an r-tightening with $r = 75$ for an image consisting of 1050 by 825 pixels. The morphological operations required three seconds while flow required four seconds over a total of six levels. The sixth level is necessary in this example because the radius value is larger than in the previous example, requiring

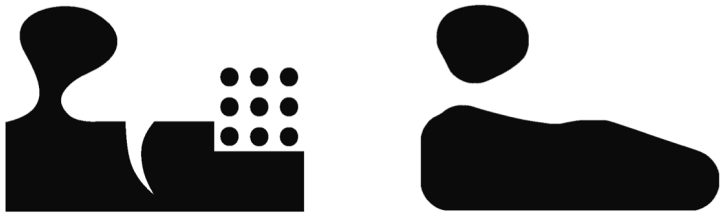


FIGURE 4.2. (Left) The input binary image consists of 1050 x 825 pixels. (Right) The tightening of the input with a radius of 75 pixels.

us to downsample an additional time. As a result, the sixth level is small, so it adds negligible execution time.

4.4. Three-dimensional implementation

4.4.1. Triangle mesh representation. In three dimensions, we avoid the visible terracing and faceting artifacts from binary volume by storing values at different voxels to obtain sub-voxel precision. Given a triangle mesh as input, we compute the signed distance from each voxel to the mesh. Our algorithm to compute distances implements a variant of Danielsson's Euclidean distance mapping [**22**]. The implementation quantifies the minimum distance from a point in a voxel to the fragments of triangles contained in a different voxel.

Using Euclidean distance mapping applied to input triangle mesh S, we extract the isosurfaces for $S \uparrow_r$ and $S \downarrow_r$ given radius r. We then compute distances for $S \uparrow_r$ and $S \downarrow_r$, yielding isosurfaces from the opening operator $S\circ_r$ and the closing operator $S\bullet_r$. We finally compute distances from $S\circ_r$ and $S\bullet_r$, requiring a total of five distance transforms.

To perform curvature flow, we begin by initializing the level set function as the distance to the boundary of the input set S. We constrain the function to be less than the distance to the boundary of $S\circ_r$ within $S\circ_r$, and we similarly constrain the function to be greater than the distance to the boundary of $S\bullet_r$ in the complement of $S\bullet_r$. As a result, the boundary of the deforming set lies within the mortar.

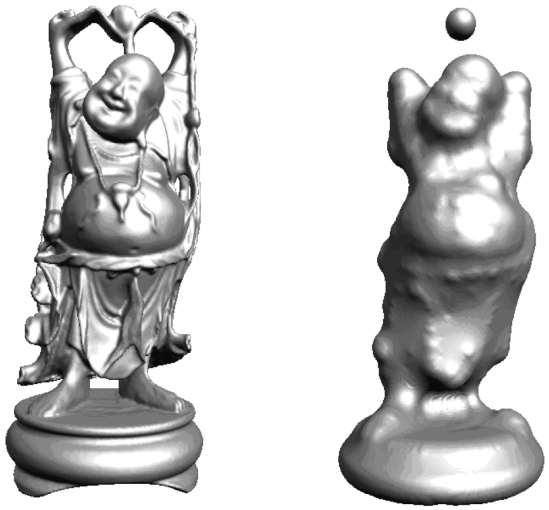


FIGURE 4.3. (Left) An input mesh consists of 829,604 triangles. (Right) The tightening of the input computed on a 128 x 128 x 279 voxel grid with a radius of 8 voxels.

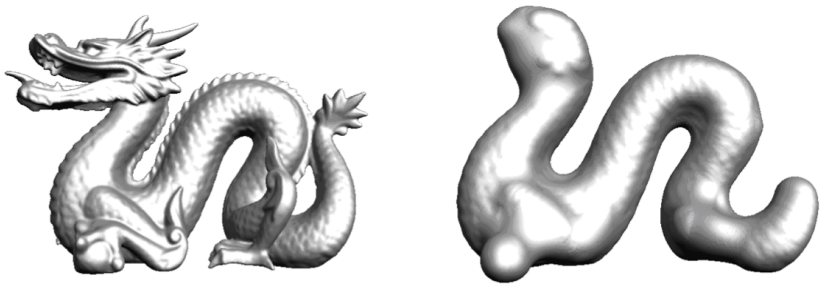


FIGURE 4.4. (Left) An input mesh consisting 1,031,894 triangles. (Right) The tightening of the input computed on a 211 x 145 x 287 voxel gird with a radius of 12 voxels.

4.4.2. Results. To the left, Figure 4.3 shows an input mesh consisting of 829,604 triangles contained in a grid composed of 128 x 128 x 279 voxels. To the right, the figure shows the tightening of the input with a radius of 8 voxels. The distance transforms individually require from 1 hr 16 min 32 sec to 1 hr 24 min 49 sec, amounting a total of 6 hr 47 min 28 sec. We additionally perform 300 iterations in a total of 17 min 7 sec.

Figure 4.4 shows to the left a second input mesh consisting of 1,031,984 triangles within a 211 x 145 x 287 voxel grid. To the right, we show a tightening of the mesh with a radius of 12 voxels. The time required by distance transforms ranges from 2 hr

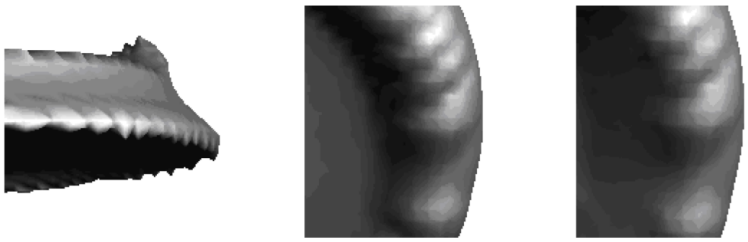


FIGURE 4.5. (Left) A sawtooth pattern that appears along the edge of the Buddha model's base when it is shrunk by a radius of 4 voxels on a 67 x 67 x 142 grid. (Center) The model on the left grown by 4 voxels, which shows a bumpy edge. (Right) Tightening by a radius of 4 voxels retains the bumpy edge generated by opening and closing.

11 min 52 sec to 2 hr 49 min 26 sec, amounting to 13 hr 6 min 5 sec. Three hundred iterations of flow then require 32 min 17 sec.

4.4.3. Artifacts. Isosurfaces extracted from distance transforms on voxel grids exhibit sharp and thin features, particularly when the grid is coarse. Tightening then reproduces such artifacts. Figure 4.5 shows a sawtooth pattern on the edge of the base of the Buddha model S, which $S \downarrow_4$ shrinks by a radius of 4 voxels on a 67 x 67 x 142 voxel grid. Growing with $\uparrow_4$ to yield the opening $S\circ_4$, we obtain a bumpy edge, which we retain in the tightening. We speculate that we may be able to remove artifacts by eliminating voxel grids and exclusively relying on triangle meshes for distance transforms and curvature flow.

4.5. Conclusion

We define r-tightening, which like opening, closing, and mason produces output sets that contain the core and exclude the anticore. However, r-tightening is distinguished by both its mean curvature bound and its locally minimal boundary measure. Because r-tightening is not unique, the r-tightening's boundary organizes components of the core into different configurations. This raises issues concerning the comparative advantages of configurations that use the mason output to resemble the input set as opposed to configurations that minimize tightening's total boundary measure.

We construct r-tightening by using constrained mean curvature flow. We implement curvature flow with level sets, which in two dimensions represent data using binary images and in three dimensions use both triangle meshes and volumes composed of voxels. We accelerate level set flow by applying both narrow bands and progressively upsampled versions of downsampled input.

As defined in Chapter 3, a closed set $S \subseteq \mathbb{R}^d$ with manifold boundary is r-regular if and only if both S° and S^c are equal to unions of open balls of radius r separated by a manifold. Consequently, S is r-regular if $S = S \circ_r \bullet_r$. The outputs generated by opening, closing, mason, and tightening are approximately regular, but r-tightening's property that its mean curvature ranges from $-\frac{1}{r}$ to $\frac{1}{r}$ helps r-tightening resemble the regular principal curvatures ranging from $-\frac{1}{r}$ to $\frac{1}{r}$. Generating regular output with minimal change to the input remains an open problem.

CHAPTER 5

FAIRING, BLURRING, BLENDING

5.1. Introduction

We present prior art techniques for smoothing sharp features, such as cusps, corners, and creases. We organize techniques in three categories: fairing, blurring, and blending. Fairing minimizes an energy functional, which optimizes a surface's local properties. Blurring adjusts the amplitudes of different surface eigenfunctions, with an emphasis on lowpass filtering by attenuating higher frequencies. Blending, by contrast, uses set-theoretic operations to alter the set of balls lying in a set and the set of balls lying in its complement.

Tightening's use of constrained mean curvature flow is a form of fairing, and it broadly relates to both fairing and blurring. Mason as a result of morphological operations is instead a form of blending. In two dimensions, the tangents connecting circular arcs in the boundary of a tightening can also be constructed by blending.

5.2. Fairing

5.2.1. Energy functionals. To fair a surface, we define an energy functional in terms of a local property at every point. A local property is typically derived from principal curvatures. For instance, the mean curvature functional used in tightening is the integral of the mean principal curvatures at each point. A variety of functionals exist, which we express in terms of flow or energy. In addition to mean curvature flow, we identify Willmore flow [**7**] and minimum curvature variation flow [**49, 50**].

5.2.2. Willmore flow. We can express the Willmore energy W of a two-dimensional smooth surface S in three dimensions with mean curvature H and total curvature K in terms of Gaussian curvature as

$$W = \int_S H^2 dA - \int_S K dA$$

We equivalently express W in terms of the principal curvatures k_1 and k_2:

$$W = \frac{1}{4} \int_S (k_1 - k_2)^2 \, dA$$

Willmore energy is zero for a sphere, because the principal curvatures are equal at every point on a sphere. More generally, Willmore flow decreases the magnitude of the difference between principal curvatures. Given constraints, differences minimize once they stabilize under flow.

5.2.3. Minimum curvature variation flow. Minimum curvature variation is equal to the integral of the squared derivative principal curvatures. Denoting the directions of the two principal curvatures k_1 and k_2 as e_1 and e_2, respectively, we express the minimum curvature variation MCV for a two-dimensional surface S as follows:

$$MCV = \int_S \frac{dk_1}{de_1}^2 + \frac{dk_2}{de_2}^2 dA$$

The energy of minimum curvature variation is zero for spheres, cylinders, cones, and torii because the derivative of each principal curvature along the corresponding principal curvature line is equal to zero. Given constraints, the flow of minimum curvature variation reduces the derivatives of principal curvatures.

5.3. Blurring

5.3.1. Lowpass filtering. In signal processing, we express a function defined over a domain as a sum of sinusoidal eigenfunctions. When the domain is the unit interval from the real line, the sinusoidal eigenfunctions have integer frequency, and the amplitude for each frequency is the amplitude of the corresponding sinusoid. In filtering, we modulate the amplitudes of the sinusoids and reconstruct the output. Lowpass filtering selectively reduces the amplitudes of high-frequency sinusoids. Qualitatively, lowpass filtering suppresses high-frequency features such as sharp edges and cusps. This rounds or softens contours, which we refer to as blurring.

5.3.2. Mesh signal processing. Mesh signal processing takes a triangle mesh as input [**68, 69**]. We construct Laplace-Beltrami eigenfunctions [**48**] for the mesh by associating a Laplace-Beltrami vector with each mesh vertex. If the mesh is a regularly sampled sphere, the eigenfunctions are discrete spherical harmonics that represent higher-dimensional sinusoids. More generally, eigenfunctions reflect the topology of the mesh's graph structure and may also reflect the geometry of the mesh vertex positions.

One approach to computing mesh eigenfunctions is to store vertex positions in a square matrix reflecting the mesh's connectivity. We then solve for the mesh eigenfunctions and express them in a series of square matrices, which we scale and add to produce filtered output. Filtering depends on how we scale the eigenfunctions.

Alternatively, we can represent mesh vertices by using one row for each vertex with three columns, each for one of the three dimensions. We next use a square matrix, where each row expresses a discrete operator for a given vertex by containing a weight for each vertex in the row. We illustrate the mean curvature operator presented by Desbrun et al. [**23**], which is analogous to a Laplace-Beltrami vector:

$$-\overline{\kappa}n = \frac{1}{4A} \sum_{j \in N_1(i)} (\cot \alpha_j + \cot \beta_j) \|x_j - x_i\|$$

Symbols x_i and x_j denote the vertices with numbers i and j, respectively. Symbol $\overline{\kappa}$ denotes the mean curvature at x_i with unit normal n, so we multiply $-\overline{\kappa}$ by n to obtain the mean curvature normal at x_i. Symbol A is the sum of the areas of triangles incident on x_i. Symbol $N_1(i)$ is the set of vertex numbers adjacent to i in the mesh, and we use j to denote a vertex number in $N_1(i)$. Symbols α_j and β_j measure angles opposite the edge e_{ij} from x_i to x_j, where α_j lies on the left of e_{ij} while β_j lies on the right.

In the equation, the value $x_j - x_i$ is a vector from x_i to x_j. The value $\cot \alpha_j + \cot \beta_j$ scales $x_j - x_i$, and the sum of scaled vectors divided by $4A$ yields the value of $\frac{\nabla A}{2A}$. Then $\frac{\nabla A}{2A}$ is equal to the mean curvature normal $-\overline{\kappa}n$.

Qualitatively, the mean curvature operator defines eigenfunctions for the mesh's square matrix. Each eigenfunction has an eigenvalue, and if we adjust mesh vertex positions by applying the mesh curvature operator each eigenfunction's amplitude decreases by an amount that increases as its eigenvalue increases. Iteratively applying the mean curvature operator performs lowpass filtering, gradually attenuating higher eigenvalues.

A variety of related techniques perform diverse forms of filtering. Although lowpass filtering is common, band-stop filtering in order to attenuate frequencies over an intermediate interval has applications to reshaping surfaces [**42**]. High-pass filtering, which attenuates low frequencies, instead removes most of the volume bounded by a surface. Reduced to a spiky shape, the surface has limited applications to modeling.

5.3.3. Solid lowpass filtering. The input to solid lowpass filtering consists of binary image data composed of pixels or voxels. We apply a Gaussian blur to the input, which performs lowpass filtering. Next, we threshold the result, rendering

binary image data as the output **[34]**. The output qualitatively reflects filtering, but we do not directly produce it by scaling eigenfunctions.

We write a one-dimensional Gaussian function G for variable x with standard deviation σ and center α as

$$G(x) := \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\alpha)^2}{2\sigma^2}}$$

Similarly, we obtain the two-dimensional Gaussian function G as the product of two one-dimensional Gaussians with variables x and y with center (α, β):

$$G(x, y) := \frac{1}{2\pi\sigma^2} e^{-\frac{(x-\alpha)^2+(y-\beta)^2}{2\sigma^2}}$$

Given an input image I, we convolve $I(x, y)$ with first $G(x)$ and then $G(y)$, blurring $I(x, y)$ by $G(x, y)$. Multiplying the denominator of $G(x)$ by $\sqrt{2\pi\sigma^2}$ to obtain the denominator $2\pi\sigma^2$ of $G(x, y)$ accounts for the quadratic number of samples within a fixed distance from center (α, β) as opposed to the linear number of samples within the same distance from α. Using a similar construction, we obtain $G(x, y, z)$ by convolving $I(x, y, z)$ with $G(x)$, $G(y)$, and $G(z)$.

The blurred output image stores a finite number of pixel or voxel values. We select a threshold from the range of values, so that values below the threshold map to binary image values inside the lowpass filtered input shape while values at or above the threshold map to image values outside. The threshold value determines the number of pixels or voxels inside the output shape; the shape becomes larger as the threshold increases, and it becomes smaller as the threshold decreases. As one possibility, we can choose a threshold value that matches the output area or volume to the area or volume of the input.

Solid lowpass filtering smooths the boundary of the output shape, minimizing both sharp features like corners and cusps and thin features like branches. Filtering also induces topological changes, separating some components of the input, merging nearby components, and eliminating small components. However, solid lowpass filtering does not offer guarantees analogous to those provided by mason or tightening.

5.4. Blending

5.4.1. Medial axis trimming. Blending a shape by trimming the medial axis of its interior or complement introduces a variety of blends and fillets [**73**]. The medial axis defines a set of balls centered on medial axis points with radii that vary continuously along the medial axis [**6, 20, 53, 77, 78**]. Trimming eliminates balls with locally minimal radii. If the trimmed medial axis lies in the input shape's interior, the balls centered on the medial axis points do not include sharp features like convexities. As a result, the union of the balls replaces certain sharp features with rounds. Similarly, trimming the medial axis in the complement of the input replaces sharp features with fillets.

The union of balls centered on medial axis points in the input shape with radii greater than or equal to r yield the input's r-opening, while the complement of comparable balls in the complement of the input yield r-closing. To establish a more formal treatment of trimming, we define maximal balls from the set of open balls $\mathbb{B}^S$ contained in set S:

Definition. 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 $b \subset a$.

Letting $\mathbb{M}^S$ denote the set of maximal balls in S, we define the medial axis of S:

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

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

Where $\pi\left(b\right)$ denotes the center of ball b and $\rho\left(b\right)$ denotes the radius of b, we let $\rho^S\left(p\right)$ for point $p \in \mathrm{MA}^S$ denote the radius of the maximal ball in S with center p. The trimmed medial axis for r-opening is then

$$\mathrm{MA}^{S \circ_r} = \bigcup \left\{ p \in \mathrm{MA}^S \mid \rho^S\left(p\right) \geq r \right\}$$

Similarly, the trimmed medial axis for r-closing is

$$\mathrm{MA}^{S \bullet_r} = \left(\bigcup \left\{ p \in \mathrm{MA}^{S^c} \mid \rho^{S^c}\left(p\right) \geq r \right\} \right)^c$$

Opening and closing trim the medial axis by thresholding radius values. Other forms of trimming can eliminate sections of the medial axis defined by its graph structure or outline boundaries of pieces of the medial axis to be removed. Alternatively, we can trim the medial axis using an adjunct shape positioned in relation to the input shape. Relative blending involves an analogous approach.

5.4.2. Relative blending. Relative blending **[73]** globally blends a shape's features by performing a combination of modulated opening and closing operations. Given an input shape, a superimposed control shape determines the blending applied to the input. We manipulate a set of balls lying in the symmetric difference of the input and the control shape, so the material outside of the set of balls is either added to the input or removed from it. Given a planar input shape, removing material rounds a convex corner, while adding material fillets a concave corner. The distance from the control shape's boundary to the input boundary varies the radius used by different rounds and fillets. In some circumstances, the radius of a round or fillet may reduce to zero, so that no blend affects a particular corner. Blends exhibit more

complex geometry in higher dimensions, while the control shape modulates all of the radii in a particular blend.

Assume that the input shape $A \subseteq \mathbb{R}^d$ and the control shape $B \subseteq \mathbb{R}^d$ are both closed. The moat, denoted by symbol X, is the closure of the symmetric difference between A and B:

$$X := (A - B^\circ) \cup (B - A^\circ)$$

The set O contains all closed balls in X that intersect both the boundary of A and the boundary of B. Let β denote a closed ball, while

$$O := \{\beta \subseteq X \mid (\beta \cap \partial A \neq \emptyset) \wedge (\beta \cap \partial B \neq \emptyset)\}$$

The pad, denoted symbol P, is the union of the balls in O:

$$P := \bigcup O$$

The mean, denoted symbol M, consists of the set of points closer to the intersection of the interiors of A and B than the complement of the union of A and B.

$$M := \left\{p \in \mathbb{R}^d \mid d\left(p, A^\circ \cap B^\circ\right) \leq d\left(p, (A \cup B)^c\right)\right\}$$

The boundary C of M is the union of centers in the balls lying in O:

$$C := \bigcup \{p \in \pi\left(\beta\right) \mid \beta \in O\}$$

Consequently, C is a subset of the medial axis of the moat.

The relative blending of A with B, denoted $R_B\left(A\right)$, modifies the input shape using combinations of the mean and the pad:

$$R_B(A) := (A \cap (M \cup P)) \cup (M - P)$$

The set $A \cap (M \cup P)$ opens the input by intersecting it with the union of the mean and the pad, where the radius of each ball in the pad is determined by the spacing between the input and control shape. The set $M - P$ similarly adds fillets to the opening, where fillets consist of the subset of the mean outside of the input minus the pad.

In an implementation of relative blending, computing the mean and pad are significant tasks. One approach is to represent all of the geometric sets involved using images composed of pixels or solids composed of voxels. To compute the mean for an image, we compute a distance field consisting for each pixel as the distance from the pixel to $A^\circ \cup B^\circ$. We then compute a second distance field where each pixel is the distance to $(A \cup B)^c$. The mean consists of pixels from the first field with values smaller than values from matching pixels in the second field. By modifying the distance field, we apply similar techniques to computing the mean for a solid.

To compute the pad, we identify pixels or voxels lying on the boundary of the mean by scanning the neighbors of each sample. For each boundary sample lying inside the mean, we rasterize a ball centered on sample p with radius $d(p, A^\circ \cap B^\circ)$, while for each sample outside the mean we rasterize a ball with radius $d(p, (A \cup B)^c)$. The pad is then the union of the rasterized balls. Given the mean and the pad, computing the relative blend requires simple, linear traversals to perform intersection, union, and difference.

In principle, we can compute a representation of the medial axis of the gap between the boundary representations of A and B. Trimming the medial axis to obtain the boundary of the mean, we construct the boundary of the pad from the mean. In general, relative blending allows a variety of implementations. It also supports a

variety of applications by facilitating the user's choice of the control shape, although its core function is to construct variable-radius rounds and fillets.

5.4.3. Relative convex hull. The relative convex hull is a maximally convex set containing one input set while disjoint from another **[64, 65]**. The definition of the relative convex hull is similar to the definition of the convex hull. For every convex set $S \subseteq \mathbb{R}^d$ with two points p and q in S, the closed segment $\overline{pq}$ lies in S. By contrast, for a convex set S relative to a set $G \subseteq \mathbb{R}^d$, segment $\overline{pq}$ lies in S if and only if it is disjoint from G. The convex hull of a given set is the intersection of all the convex sets containing it, while the convex hull of some set relative to G is the intersection of all convex sets containing it relative to G. Denoting the convex hull of set S relative to set G as $\mathrm{CH}(S \mid G)$, we formally define $\mathrm{CH}(S \mid G)$ as

Definition. 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 $\mathrm{CH}(S \mid G)$, is the intersection of all sets convex relative to G containing S.

Qualitatively, a set disjoint from G has a locally convex intersection or set difference with a convex hull relative to G. The relative convex hull $\mathrm{CH}(S \mid G)$ also connects components of S such that for every pair of components an edge disjoint from G connects the pair. In effect, the relative convex hull is a constrained convex hull that resembles an unconstrained convex hull.

In two dimensions, the convex hull of the r-opening of a set S relative to the complement of the r-closing of S is an r-tightening of S. Due to the use of opening and closing, the boundary of the relative convex hull has curvature between $-\frac{1}{r}$ and $\frac{1}{r}$, where the subset of the boundary disjoint from the core and anticore consists of line

segments. In this situation, the relative convex hull is the tightening with the largest area and the smallest number of components.

In three dimensions, the relative convex hull does not guarantee an r-tightening. The unsupported boundary of an r-tightening consists of saddle points with zero mean curvature, while unsupported boundary points from a relative convex hull are locally developable, with zero Gaussian curvature because the hull is locally convex. Like mason, however, the relative convex hull provides a general set-theoretic means for constructing a boundary in the mortar separating the core and anticore.

5.5. Conclusion

We present three categories of techniques relevant to the smoothing features we obtain from mason and tightening. Fairing optimizes local surface properties, blurring attenuates high-frequency eigenfunctions, and blending manipulates the set of balls defined by the input. Fairing and blurring can be applied to mean curvature flow for use in tightening, and in two dimensions the relative convex hull in blending constructs a particular tightening. The relative convex hull also solves a problem related to mason, while aspects of mathematical morphology from blending are generally applicable to mason as well.

Collectively, fairing, blurring, and blending encompass a broad range of techniques for use in solid modeling. As tools for transforming both the interior and boundary of a shape, they provide multiple approaches to a given problem.

PART 2

HULLS, BLENDS, AND COVERS

CHAPTER 6

TIGHT HULLS

6.1. 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 6.1.) 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 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, consider

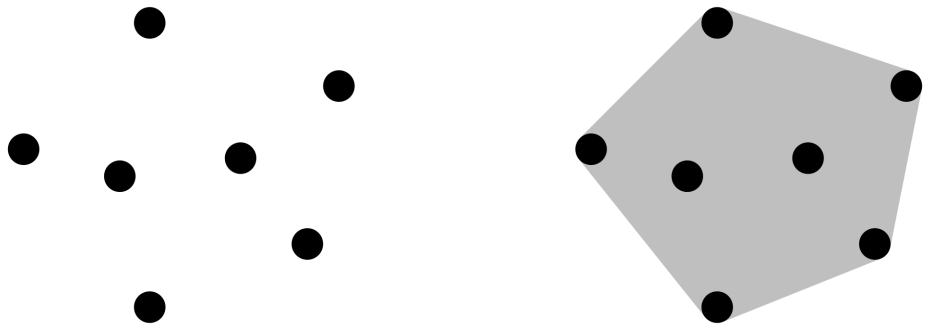


FIGURE 6.1. A set S consisting of isolated points is shown in black on the left. The convex hull of S is shown in gray on the right, with S superimposed.

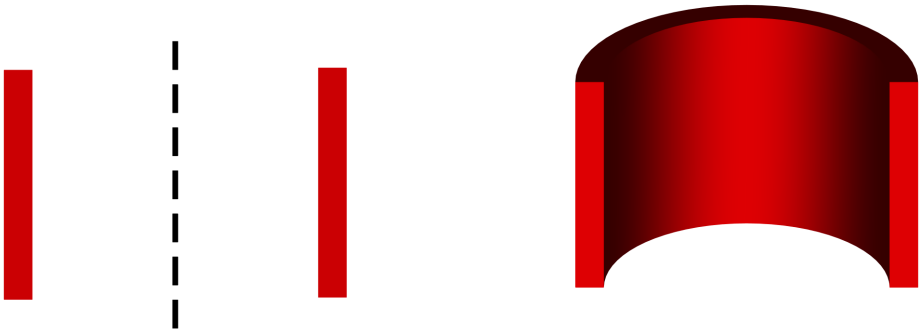


FIGURE 6.2. On the left we show a cross section of a hollow cylinder, which consists of two thin rectangles shown in red. If we rotate the rectangles about the dashed axis, they sweep the hollow cylinder, shown in a cutaway view on the right.

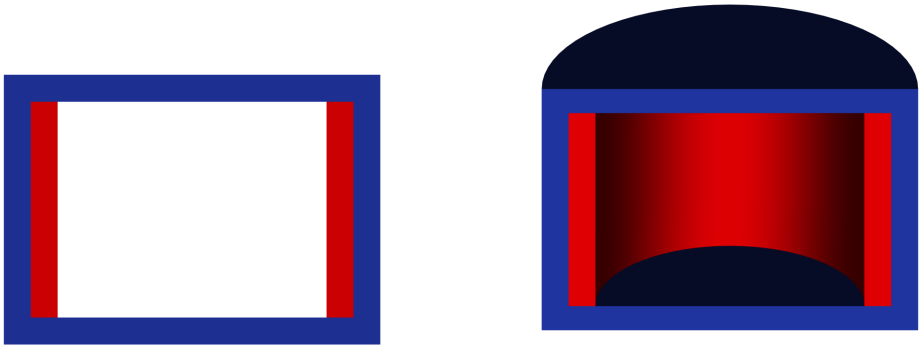


FIGURE 6.3. On the left, we show a cross section through the hollow cylinder's convex hull, with the hollow cylinder in red and the hull boundary in blue. We thicken the convex hull's boundary and offset it from the cylinder to aid visualization, but the curved portion of boundary of the hull actually lies on the boundary of the cylinder. On the right, we show a cutaway view of the cylinder and its hull.

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 6.2.) When wrapped in its convex hull, the hollow cylinder resembles a double-headed drum (Figure 6.3.) 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

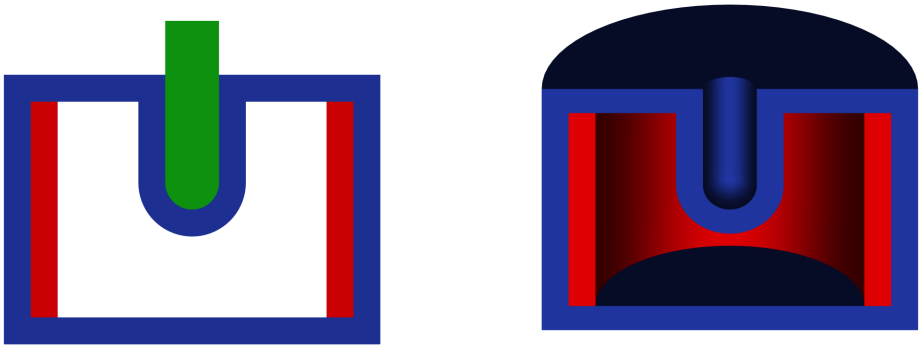


FIGURE 6.4. On the left, we show a cross section of a possible hull in blue whose boundary separates the hollow red cylinder from the interior of a green set corresponding to a finger pressing into the cylinder's convex hull. On the right, we show a cutaway of the three-dimensional hull. In both views, the hull surrounds the finger as if the hull boundary were fluid. In this respect, it behaves like a relative convex hull (Definition 6.2.)

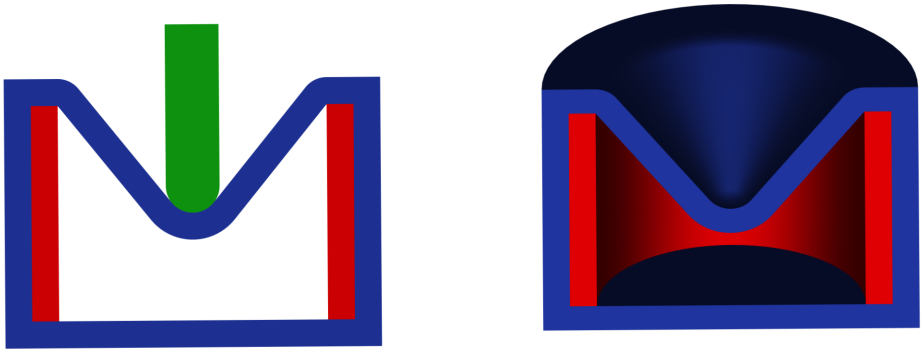


FIGURE 6.5. On the right, we show a cross section of a possible hull in blue whose boundary separates the red cylinder from a green set corresponding to a finger pressing into the cylinder's convex hull. On the right, we show a cutaway of the three-dimensional hull. In both views, the hull boundary behaves like a membrane in tension, stretched by the finger's pressure. In this respect, the hull behaves like a tight hull (Subsection 6.2.3.)

hull (Figure 6.4.) 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 6.5.)

The relative convex hull, originally defined by Sklansky and Kibler **[64]**, has a fluid-like boundary (Figure 6.4,) while the tight hull we introduce in this thesis has a boundary that behaves like a membrane in tension (Figure 6.5.) 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 6.21,) but both sets have boundaries that behave like rubber bands. In higher dimensions, relative convex hulls generally do not. We argue that rubber bands provide useful properties, making tight hulls more suitable than relative convex hulls for some applications.

6.2. Comparison of hull definitions

6.2.1. 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 **[72]**:

Definition 6.1. 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 6.6.)

The *convex hull* of $S \subseteq \mathbb{R}^d$ is the intersection of all convex sets containing S (Figure 6.7.)

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.

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

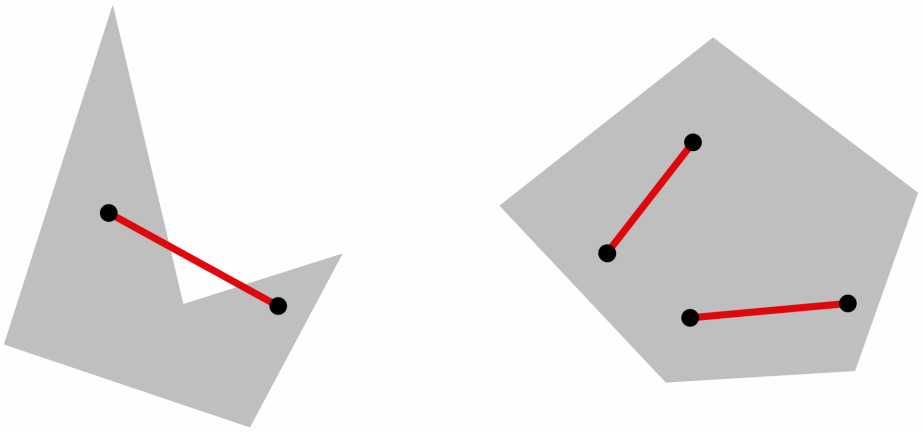


FIGURE 6.6. The gray polygon on the left is not convex because it contains the two black points but not the red segment connecting them. The polygon on the right is convex because it contains every segment that connects a pair of its points.

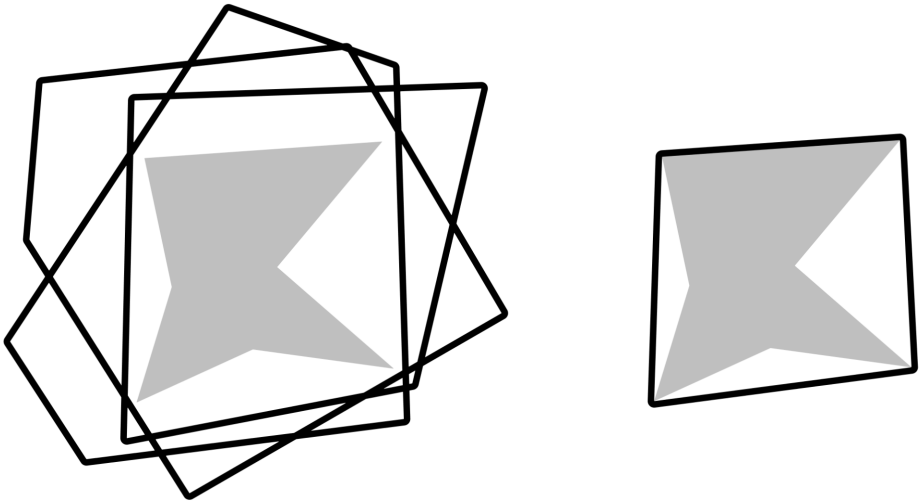


FIGURE 6.7. On the left we show a set S in gray and the boundaries of several of the convex sets containing S in black. On the right, we show the boundary of the convex hull of S in black. The convex hull of S is the intersection of the convex sets containing S.

filtering steps below, we provide a name in parentheses for the property guaranteed by each step:

(1) (Containment) Start with all subsets of Euclidean space that contain S (Figure 6.8, left.)

(2) (Convexity) Extract the convex sets (Figure 6.8, middle.)

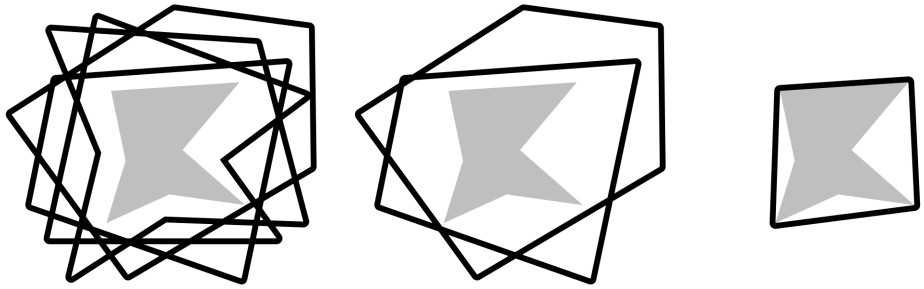


FIGURE 6.8. On the left, we show in black the boundaries of several of the sets containing the set S shown in gray. In the middle, we show the boundaries of convex sets containing S, and on the right we show the boundary of the smallest convex set containing S, which is the convex hull of S.

(3) (Smallness) The convex hull of S is the smallest of the remaining sets when they are ordered by set inclusion (Figure 6.8, 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.

6.2.2. 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 [**64**] and Sloboda and Zatko [**65**] to facilitate our later analysis of symmetry with respect to set complement (Section 6.4.)

Definition 6.2. 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 6.9.)

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 6.10.)

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

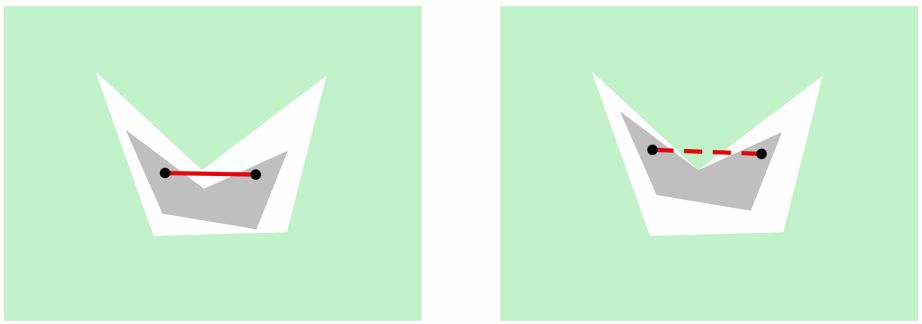


FIGURE 6.9. The set S shown in gray on the left is not convex relative to the green set G because the red line segment connecting the two black points is not contained in S despite the fact that it is disjoint from G. The set shown in gray on the right, by contrast, is convex relative to G.

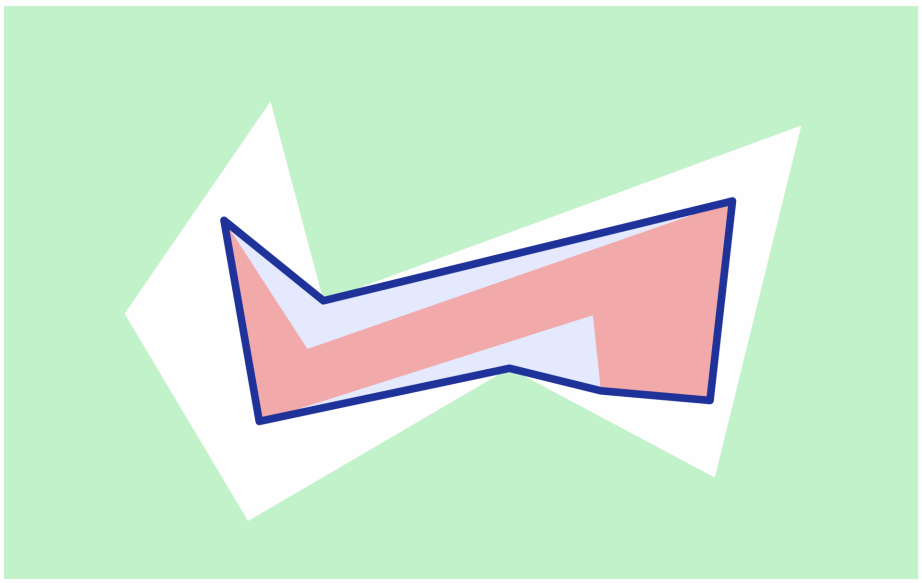


FIGURE 6.10. This image shows the boundary of the convex hull of R relative to G in dark blue, where R and G are shaded red and green, respectively. The hull contains R, and the material it adds to R is shaded light blue.

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 **[72]** 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 6.1. When discussing convex structures, it is typical to use the terms "convex" and "convex hull" in more inclusive senses, while the term "Euclidean convex hull" refers to the convex hull specified by Definition 6.1. We handle topological terms in a similar manner. We use a term like "closure" to refer



FIGURE 6.11. 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.

to the topological closure on $\mathbb{R}^d$ defined by the Euclidean metric, but in the context of a discussion of closure structures it refers to any closure operator.

6.2.3. 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 **[18, 45, 76]**. 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 6.11,) while any smooth surface bounding a convex set is a tight embedding of a topological sphere (Figure 6.12.)

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 equal to the difference between 2π and the sum of the angles formed by consecutive edges incident on the vertex (Figure 6.13.)



FIGURE 6.12. 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.

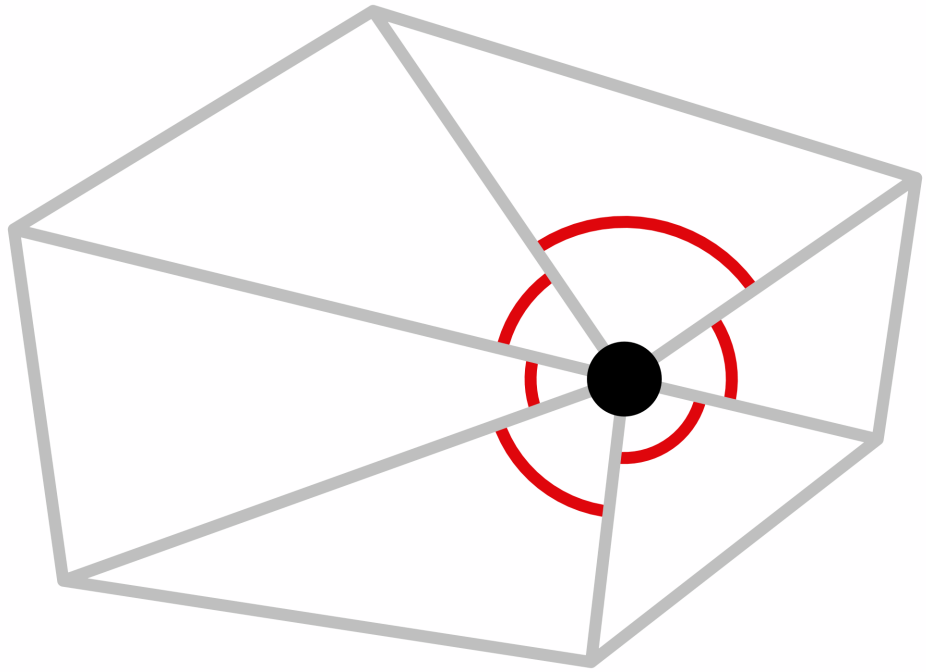


FIGURE 6.13. A vertex of a triangle mesh is shown in black with the triangles incident on it outlined in gray. The Gaussian curvature at the black vertex is 2π minus the sum of the angles subtended by the red arcs.

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 and the length of a geodesic circle on the surface centered at the point when the two circles' radii are small and equal (Figure 6.14.) At a saddle, the length of the geodesic circle

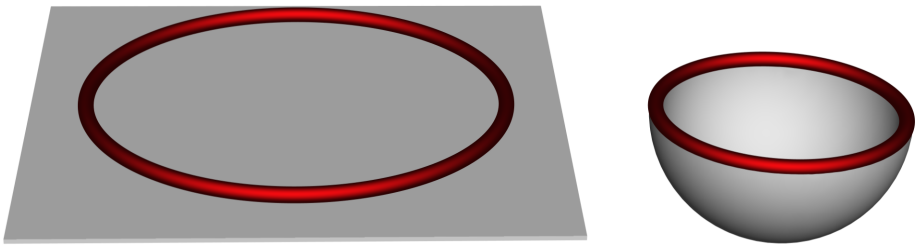


FIGURE 6.14. The surface on the left is planar, while the surface on the right is half of a unit sphere. The curves shown in red are geodesic circles of radius $\frac{\pi}{2}$. Because the surface on the right is curved, the geodesic circle on it has length 2π, while the circle on the left has length 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 $\frac{3}{\pi r^3}$ [**5**].

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 can 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. A convex vertex may have negative Gaussian curvature, while 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. 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 6.15.)

In three dimensions, we can unify smooth and polyhedral tightness [**8**] using the two-piece property [**3**]. A subset of d-dimensional Euclidean space has the two-piece

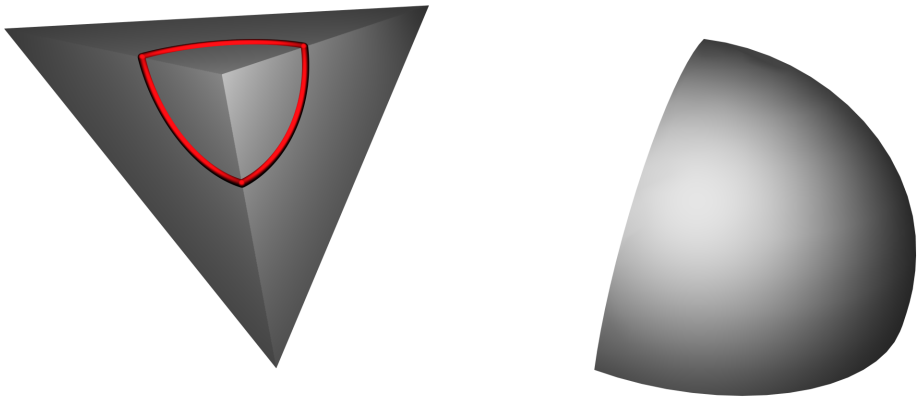


FIGURE 6.15. On the left we show a corner of a cube. No matter how we triangulate the cube's boundary, the sum of the angles incident on the vertex at one of its corners is equal to $\frac{3\pi}{2}$. Under our definition of polyhedral Gaussian curvature, the Gaussian curvature at a corner is then $2\pi - \frac{3\pi}{2} = \frac{\pi}{2}$. 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 $\frac{1}{r^2}$.

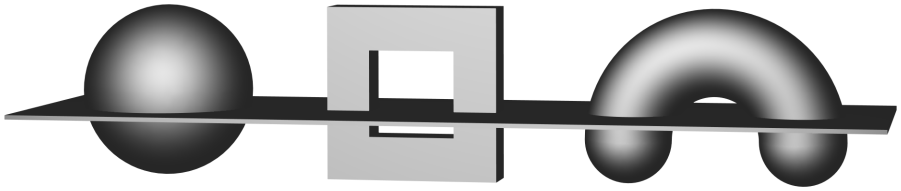


FIGURE 6.16. A sphere and a polygonal torus have the two-piece property, as illustrated on the left and in the center. The set on the right is cut into three connected components by a plane, so it is not tight.

property if and only if its intersection with a halfspace has no more than one connected component; 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 6.16.)

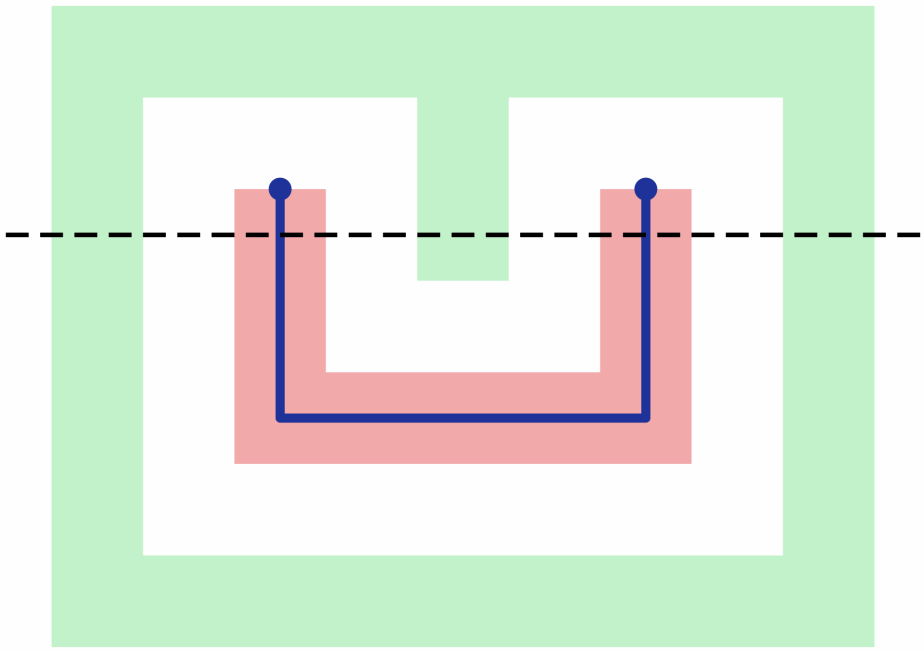


FIGURE 6.17. In the two-dimensional example above, no set that contains R and excludes G has the two-piece property. A set that contains R must contain the two points on R indicated by blue dots, but any path connecting those two points must cross the dashed black line, implying the points are in two 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.

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 that extends the definition of total absolute curvature to a variety of subsets of Euclidean space [**46**], which we formally specialize to polyhedra (Sections A.1-A.3.) 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 6.17.) This prevents us from directly applying the two-piece property to our definition of a tight hull. Moreover, in higher dimensions slack 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 [**4**]. We

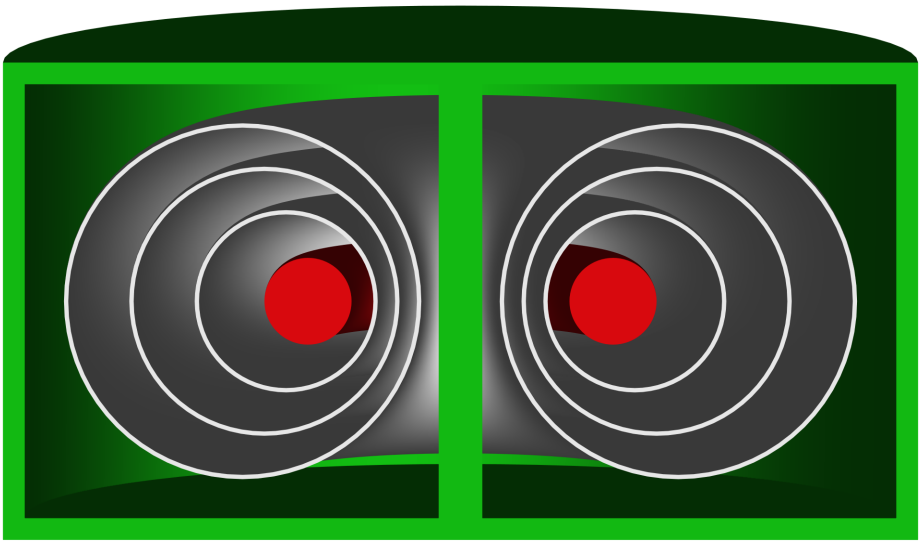


FIGURE 6.18. In this figure, R is the red torus and the complement of G is a punctured cylinder. All of the gray torus boundaries shown separate R and G with minimal slack. Because the radii of the circular torus cross sections can vary continuously, there are uncountably many torii of this kind.

therefore broadly characterize tight sets as slack minimizers. This characterization remains meaningful regardless of spatial dimension or the constraints on the slack minimization.

We conjecture that under modest assumptions there is always a slack minimizer that contains R and excludes G (Subsection A.7.2,) and in general there are uncountably many (Figure 6.18.) 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 (Section A.4.) 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.

Support, like convexity, is a value in $[0, 1]$ (Section A.5.) The support of a subset of the boundary of S from T is the integral of the support of its outward-pointing

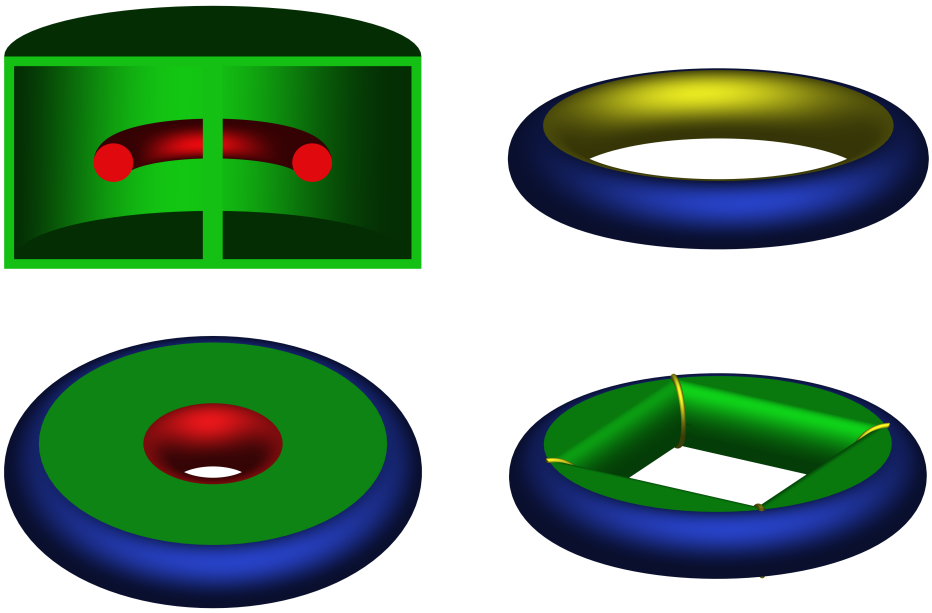


FIGURE 6.19. On the top left, we show sets R and G identical to the R and G in Figure 6.18. The tight hull of R relative to G is equal to R. We show the tight hull on the top right, with the boundary colored to indicate the amount of support at each point. The blue part of the boundary on the outer wall 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.

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 6.19, Subsection A.5.3.) If the support of a subset of a candidate hull's boundary is zero, we refer to the subset as unsupported.

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

Definition 6.3. 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

(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.

We provide a formal definition in Section A.6. We also prove that the tight hull generalizes the convex hull under limited conditions: the convex hull has minimal slack because it is convex, and all of its slack is supported (Figure 6.20, Subsection A.7.4.) 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.

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

We conjecture that a tight hull always exists, but its interior may not be unique (Figure 6.21.) 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.

6.2.4. 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

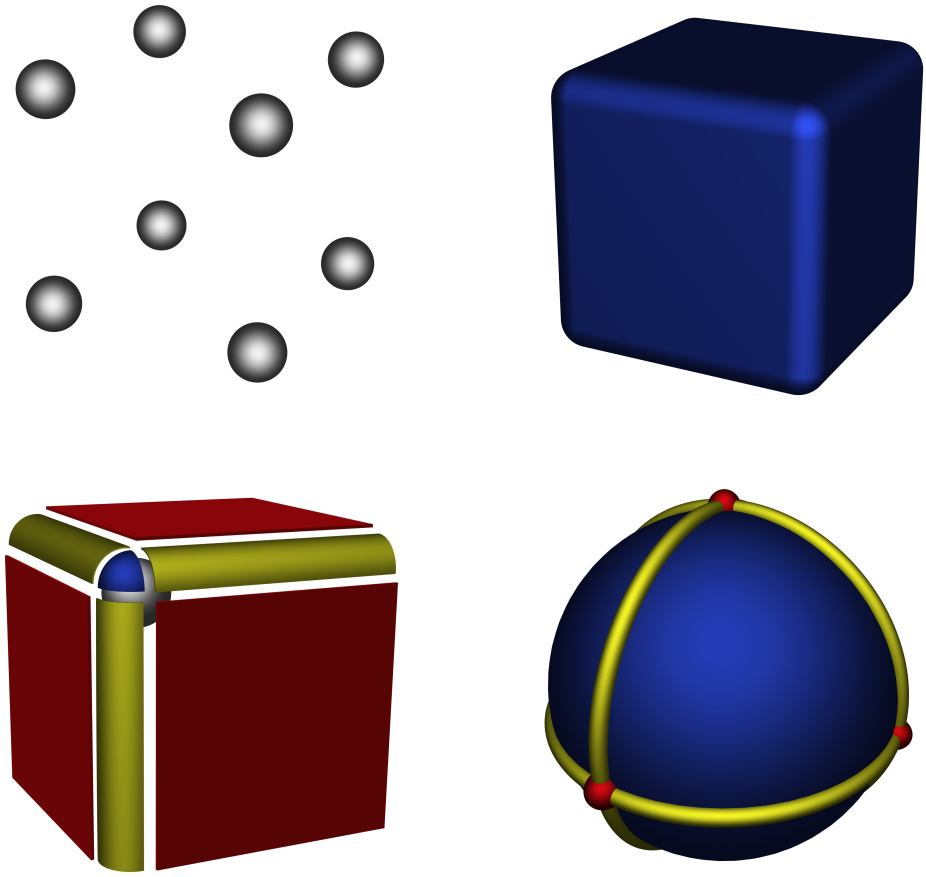


FIGURE 6.20. At the top left, we show eight small spheres, each centered on the corner of a unit cube. At the top right, we show the convex hull of the spheres in blue. At the bottom left, we show an exploded view of one of the hull's corners, with the hull decomposed into a blue spherical cap, yellow cylindrical edges, and planar red faces. On the bottom right, we 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.

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.

An intersection of tight sets, however, is not always tight (Figure 6.22.) Consequently,

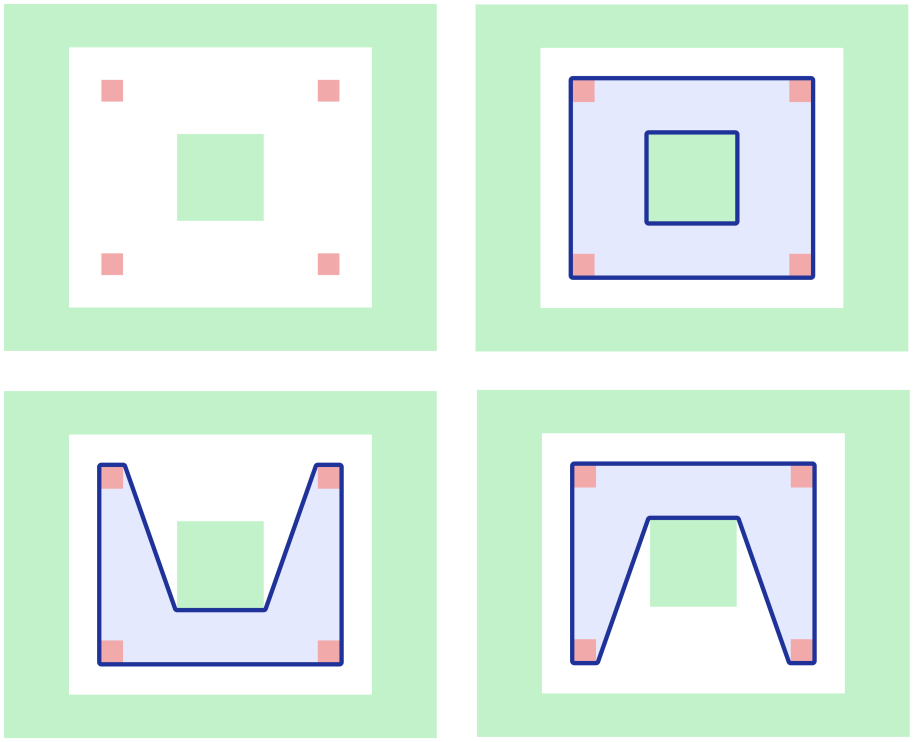


FIGURE 6.21. We show the sets R and G in the upper left. The set R consists of four red squares, while the complement of G is a rectangle with a square hole. We show the convex hull of R relative to G in the upper right; note that the hull does not minimize slack, because the sum of the angles around 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.

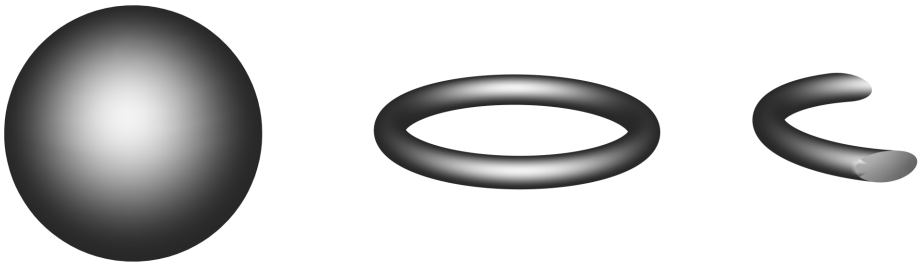


FIGURE 6.22. The sphere on the left and torus in the center are tight, but their intersection on the right is not.

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 or 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.

6.3. 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 mirror symmetry about line l (Figure 6.23.) 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 6.24,) and 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 6.25.) The set R' swept by R is the union of a hollow cylinder and a narrow cylinder with a wide, thin cylindrical base (Figure 6.26.) Similarly, G' is the union of a hollow cylinder with the unbounded complement of a cylinder (Figure 6.27.)

Figure 6.28 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

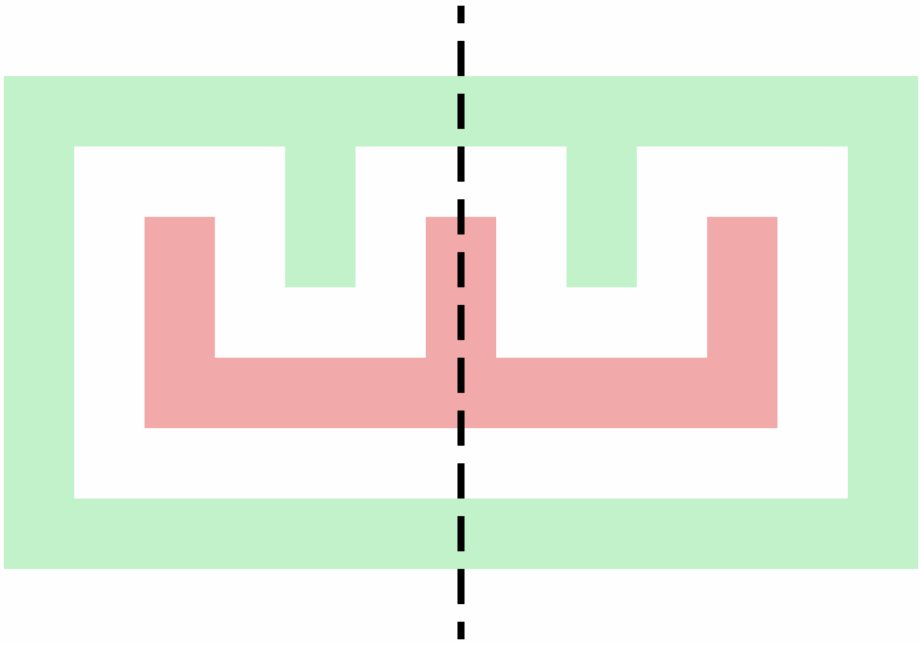


FIGURE 6.23. The sets R and G are shown in red and green, respectively. Both R and G^c are simple polygons symmetric about the dashed black line l.

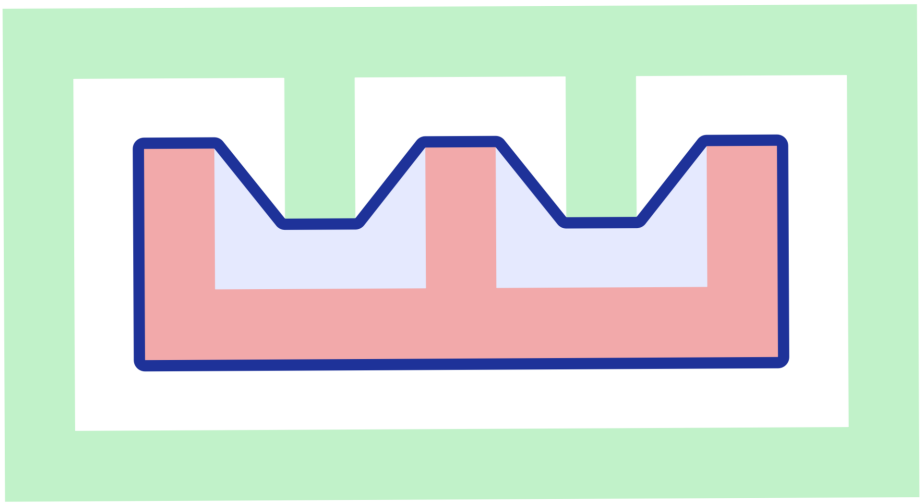


FIGURE 6.24. For the red set R and green set G shown, the convex hull of R relative to G and the tight hull of R relative to G are identical. The hull boundary is shown in dark blue, and the material the hull adds to R is shaded light blue.

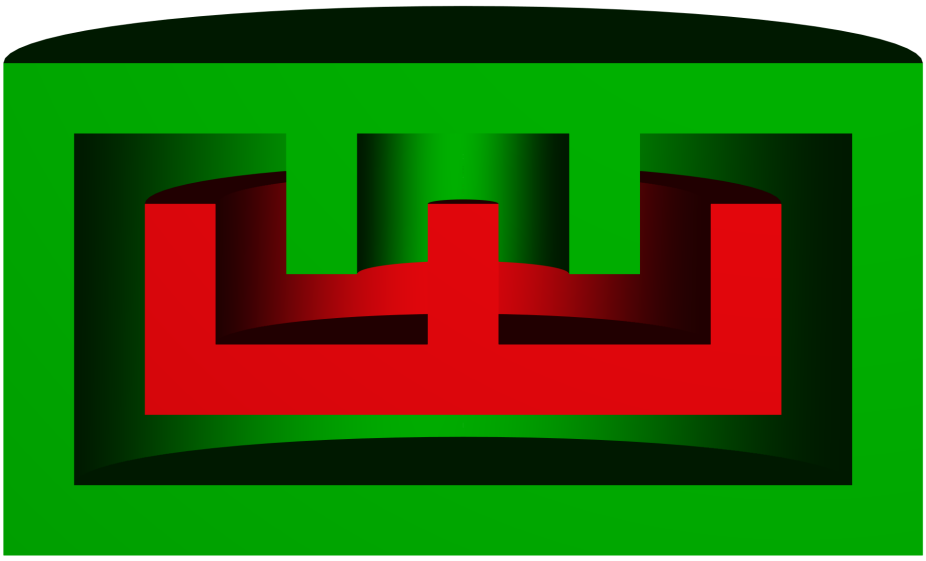


FIGURE 6.25. The red set R' and green set G' produced by rotating the planar sets R and G in Figure 6.23 about their line of symmetry l.

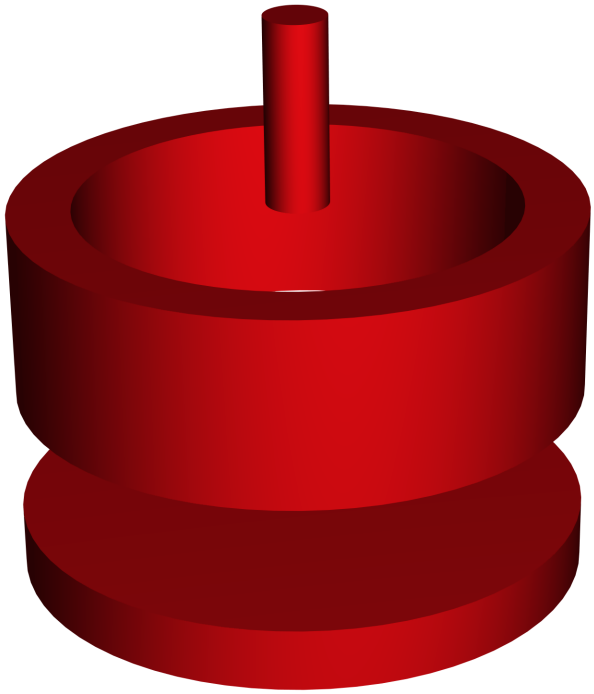


FIGURE 6.26. The set R' produced by rotating the set R shown in Figure 6.23 consists of a hollow cylinder on a wide cylindrical base with a narrow cylinder along their common axis. We show the three cylindrical pieces separately in this exploded view.

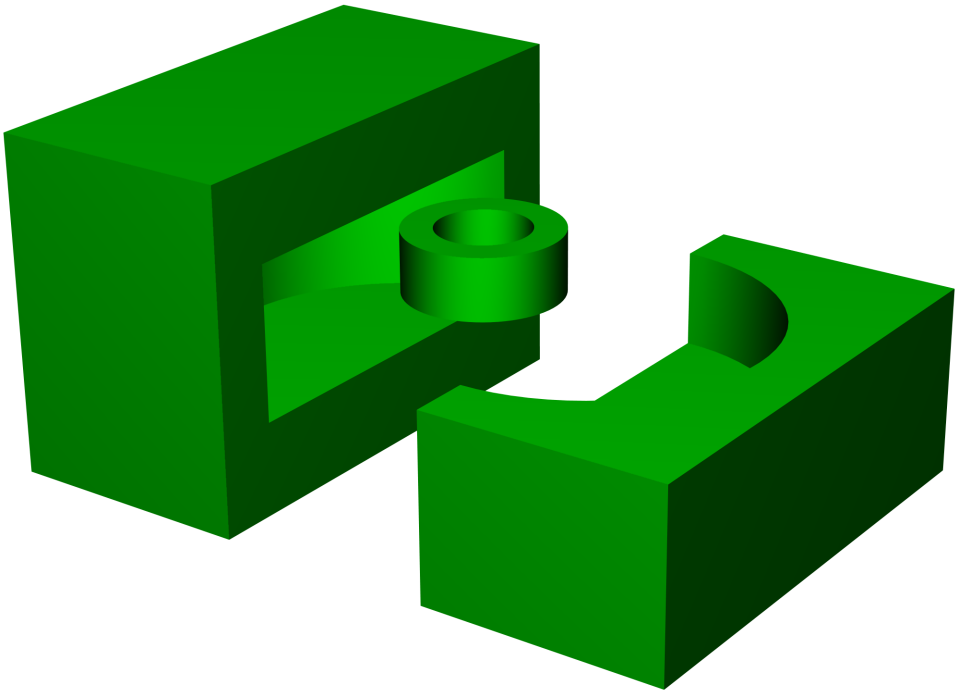


FIGURE 6.27. The set G' produced by rotating the set G shown in Figure 6.23 consists of the union of a hollow cylinder with the unbounded complement of a larger cylinder. We show an exploded cutaway of G' to make the hollow cylinder visible.

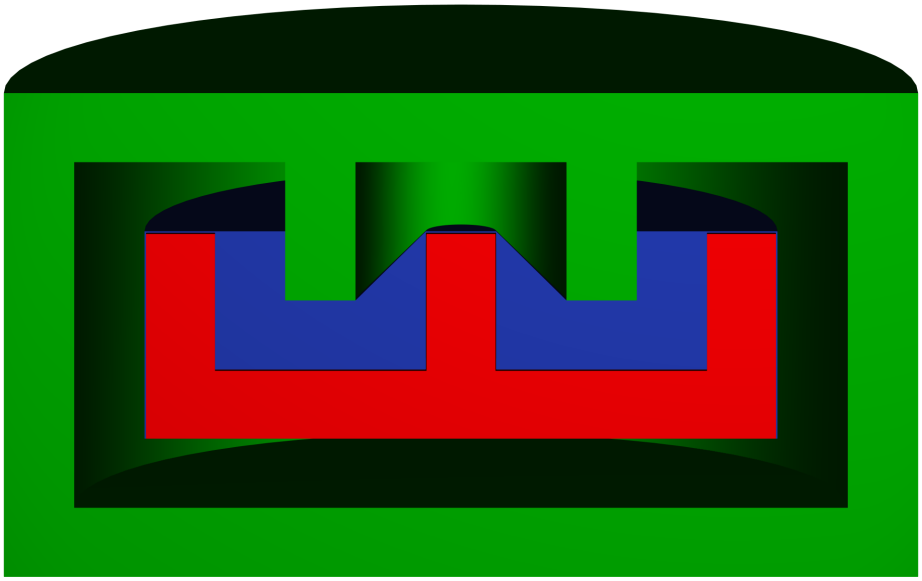


FIGURE 6.28. This figure shows a cutaway of the convex hull of R' relative to G', where R' and G' are the red and green sets shown in Figure 6.25. The relative convex hull contains R', and the material the hull adds to R' is shown in blue.

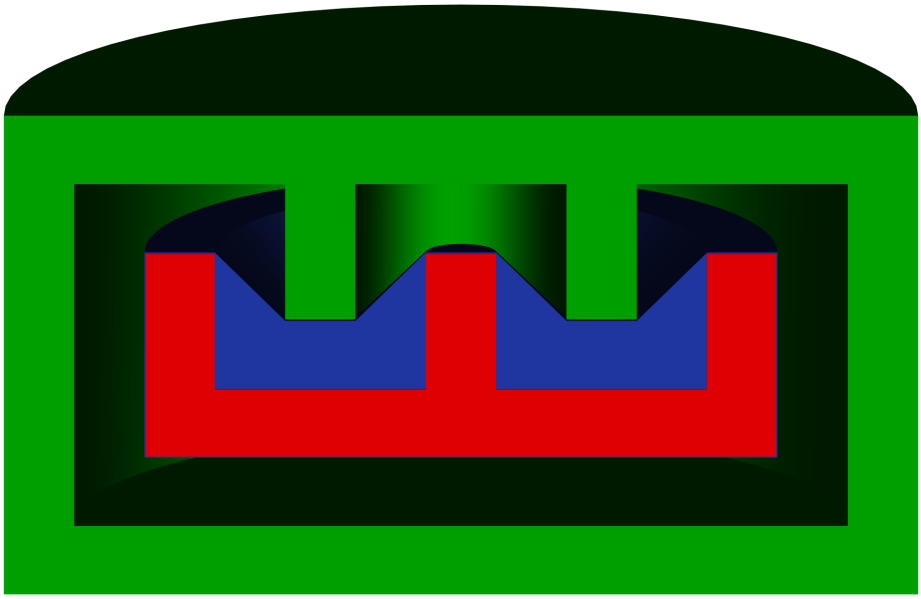


FIGURE 6.29. This image shows a cutaway of the tight hull of R' relative to G' in blue, where R' and G' are the red and green sets shown in Figure 6.25. The tight hull contains R', and a cross section of R' is visible at the cutting plane.

boundary is locally convex, which is a property of all relative convex hulls. 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 the hull. The convex hull of R relative to G is then not convex relative to G, which is a contradiction.

Figure 6.29 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 6.30.) The two angled segments in the relative convex hull section sweep out the conical portion of the hull's

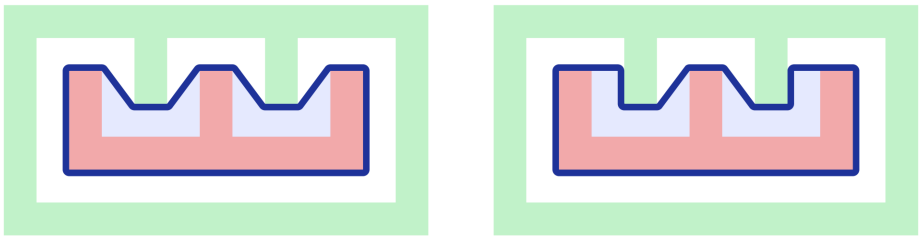


FIGURE 6.30. On the left, we show a cross section of the three-dimensional tight hull depicted in Figure 6.29. The section is identical to both the tight hull and convex hull of the planar sets R relative to G shown in Figure 6.24, so that three-dimensional tight hull is equal to the sweep of the two-dimensional 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 6.28. 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.

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 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

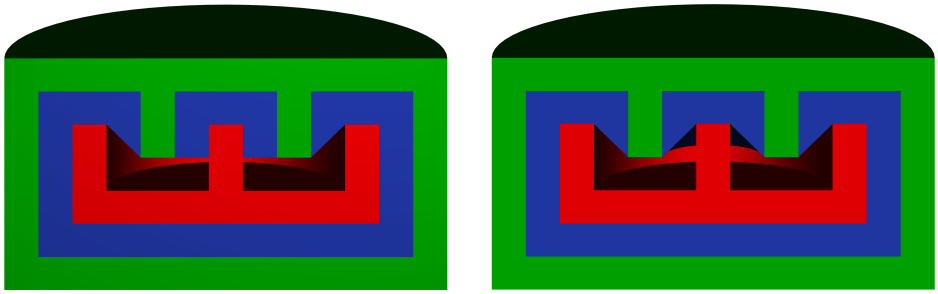


FIGURE 6.31. On the left, we show a cutaway of the convex hull of G' relative to R'. As before, the hull's unsupported boundary is locally convex. On the right, we show the corresponding tight hull. The unsupported boundary of the tight hull has both locally convex and locally concave components.

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.

6.4. 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 6.31.) 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 6.32.) 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. 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

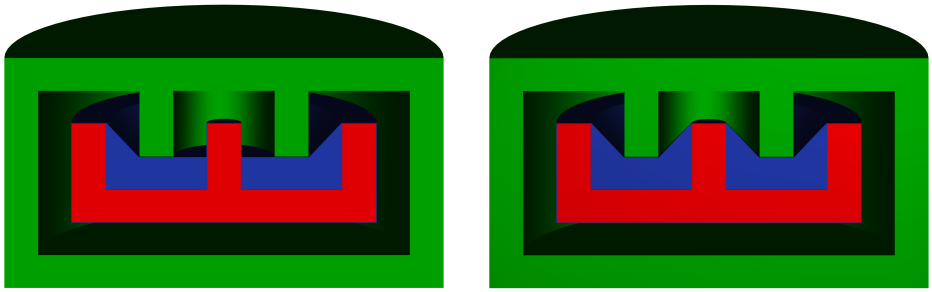


FIGURE 6.32. The set in blue on the left is the complement of the convex hull of G' relative R'. It appears substantially different from the convex hull of R' relative to G' shown in Figure 6.28. In particular, its unsupported boundary has nonpositive curvature, while the curvature of the unsupported boundary of the hull in Figure 6.28 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 6.29.

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

6.5. 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 6.33.) From the top, the hull boundary resembles a tablecloth draped over the balls, but without any folds near the

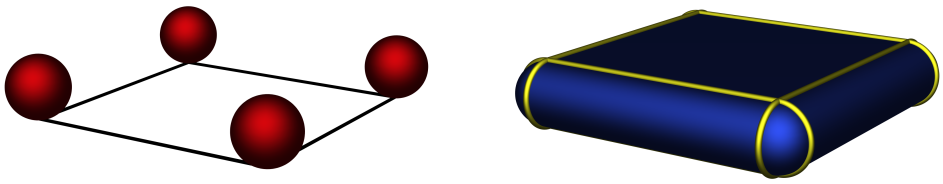


FIGURE 6.33. On the left we show four red balls of equal radius resting on the corners of a black square. The convex hull of the spheres, shown in blue on the right, consists of two squares, four cylindrical patches, and four spherical caps. We highlight the patch boundaries with yellow curves.

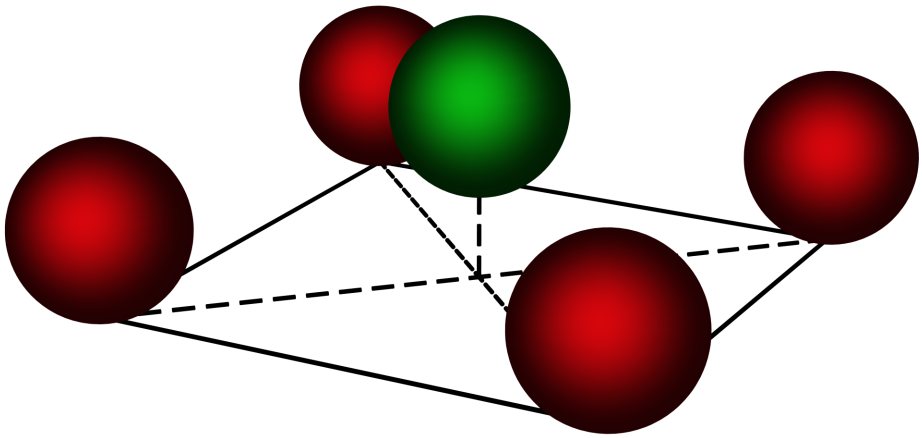


FIGURE 6.34. The set R consists of the four red balls of radius r arranged at the corners of a square. The set G is the green ball with its center at a distance of $2r$ above the square's center.

spherical caps. As in our drum example in Subsection 6.1, 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 6.34.) Figure 6.35 shows the convex hull of R relative to G, while Figure 6.36 shows the tight hull of R relative to G. As in our previous examples, the relative convex hull behaves like a fluid, surrounding the green ball. The green 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.

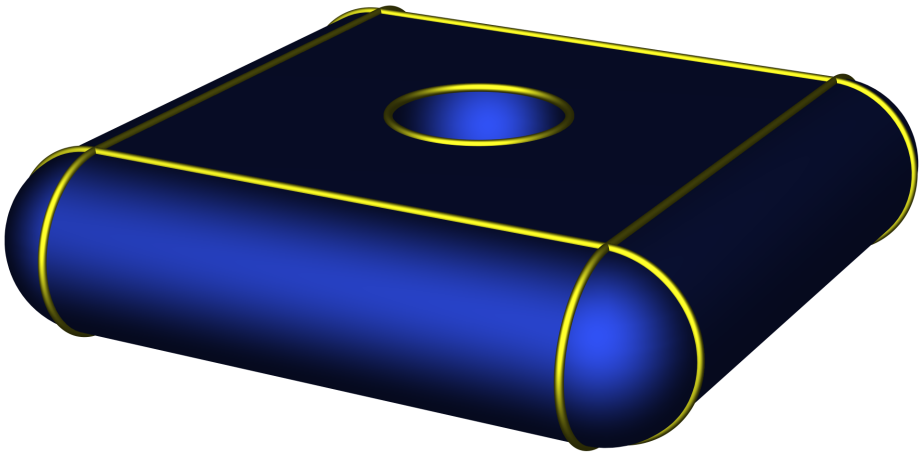


FIGURE 6.35. We show the convex hull of R relative to G in blue. The hull boundary consists of round cylinders, spherical caps, and a square with a circular hole. We highlight the patch boundaries in yellow.

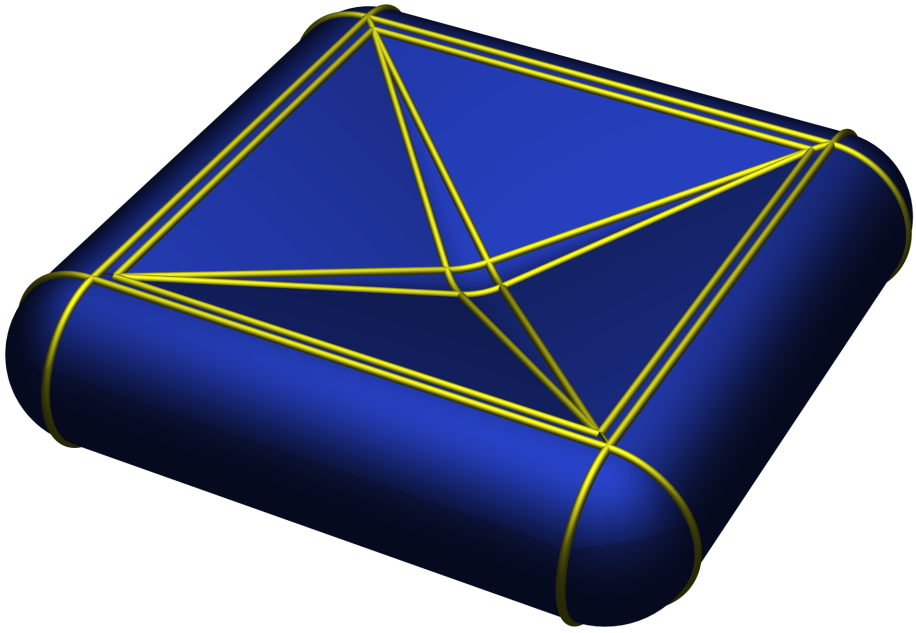


FIGURE 6.36. We show the tight hull of R relative to G in blue. The hull boundary consists of round cylinders, elliptical cylinders, spherical caps, and triangles. We highlight patch boundaries in yellow.

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

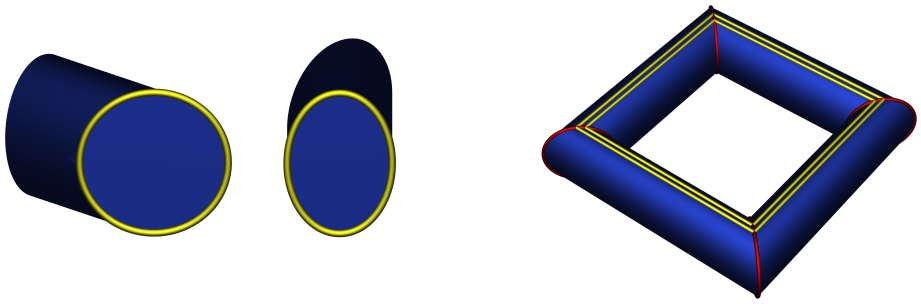


FIGURE 6.37. On the left, we illustrate how an elliptic cylinder is swept by a circle as it moves along a line that forms an angle of $\frac{\pi}{2}$ with the plane containing the circle. The cross section orthogonal to the sweep direction is elliptical. On the right, we show the elliptic cylinders swept 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.

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 the circle. (Figure 6.37.) To explain the elliptic cylinders' presence, consider the similar, but incorrect, tight hull shown in Figure 6.38. We can construct this hull by first taking the union of the convex hulls of pairs of nondiagonal red balls, forming a 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 that is suitable for nonsmooth surfaces. According to the Gauss-Bonnet theorem [**44, 67**], the integral of Gaussian curvature over a watertight surface

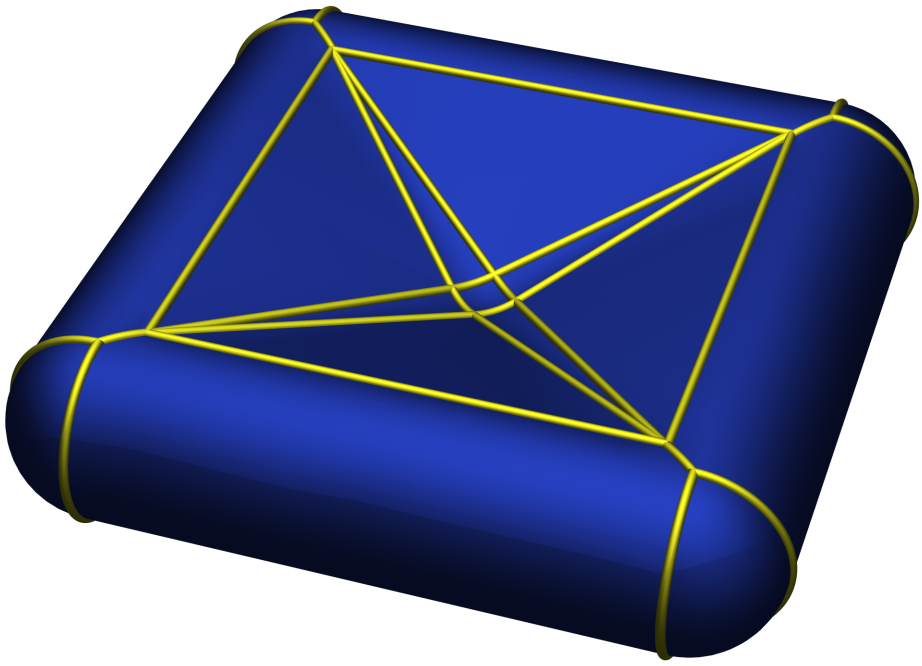


FIGURE 6.38. This hull consists of patches from round cylinders, cones, planes, and spheres. We outline the patches in yellow. In contrast to the tight hull boundary in Figure 6.36, the boundary from this hull does not contain patches from elliptical cylinders.

is determined solely by its topological type. The portions of 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.

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

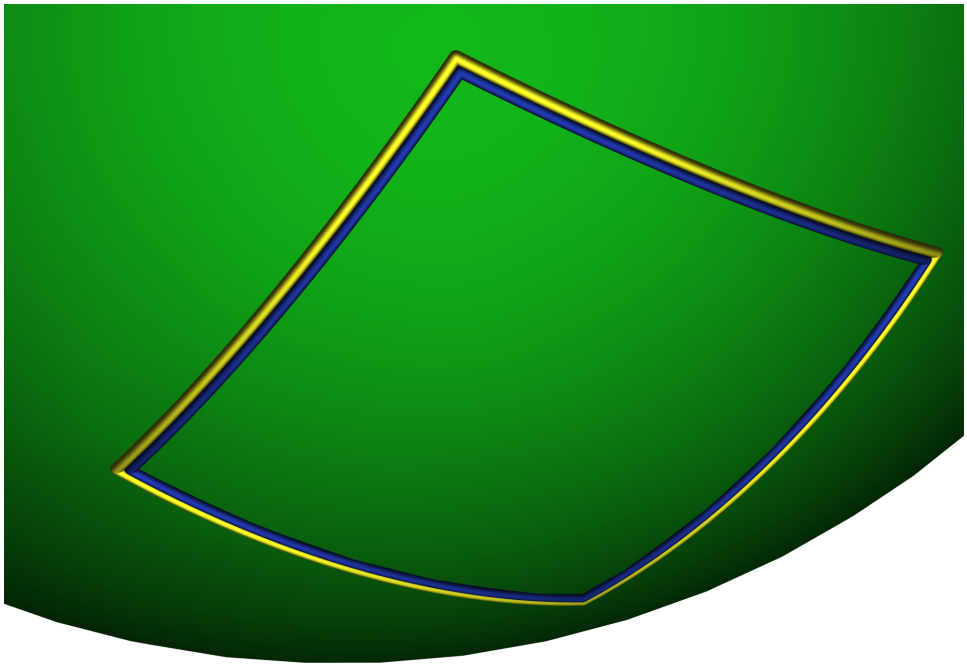


FIGURE 6.39. We show the green ball with the patch it contributes to the tight hull outlined in blue and the patch it contributes to the incorrect hull outlined in yellow. Although the patches are similar, the patch in the incorrect hull is slightly larger.

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 6.39,) and because spheres have constant Gaussian 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 6.40.)

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 sphere (Figure 6.41.) This set has the same slack as the tight hull of R relative to G, but less of its

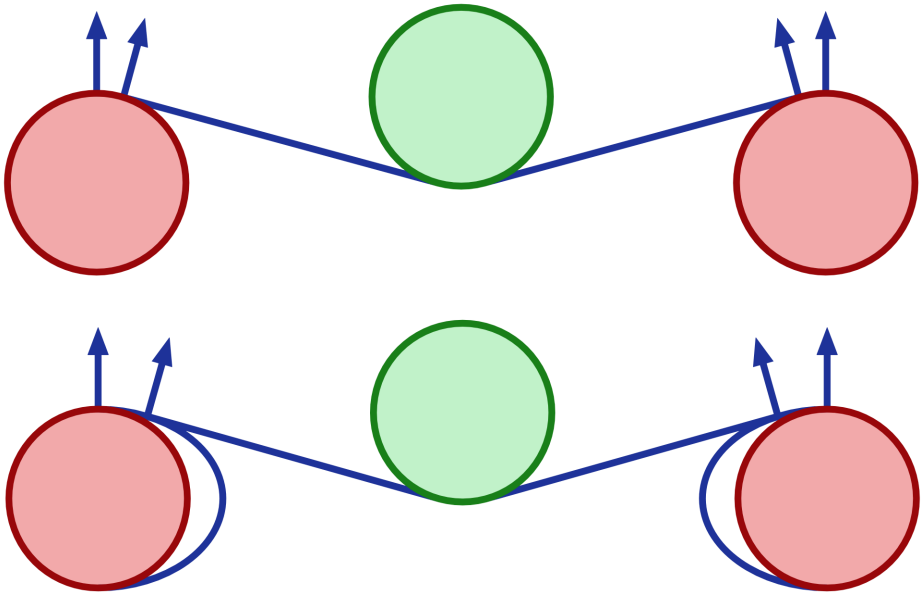


FIGURE 6.40. At the top, we show a section of the tight hull through the centers of the green sphere and two opposing red spheres. The two blue arrows indicate the change in normal direction over the circular arc lying on the edge of the elliptical torus patch in the tight hull boundary. We show the 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 section, corresponding to the smaller area of the patch in the tight hull from the green ball (Figure 6.39.)

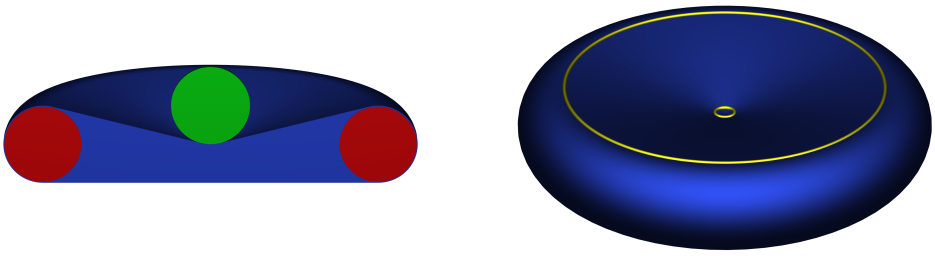


FIGURE 6.41. On the left, we show a cutaway view of the tight hull of the red torus R' relative to the green ball G'. In the image, R' is red, G' is green, and the hull is blue. On the right, we outline the toroidal, conical, and spherical patches in the hull boundary in yellow. There is an additional patch in the shape of a disk on the hull's bottom.

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.

CHAPTER 7

TIGHT BLENDS

7.1. Tight blends bound 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 have rubber band-like boundaries that press against the convex portions of the boundaries of R and G (Figure 7.1.) 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 7.2.) The boundary of

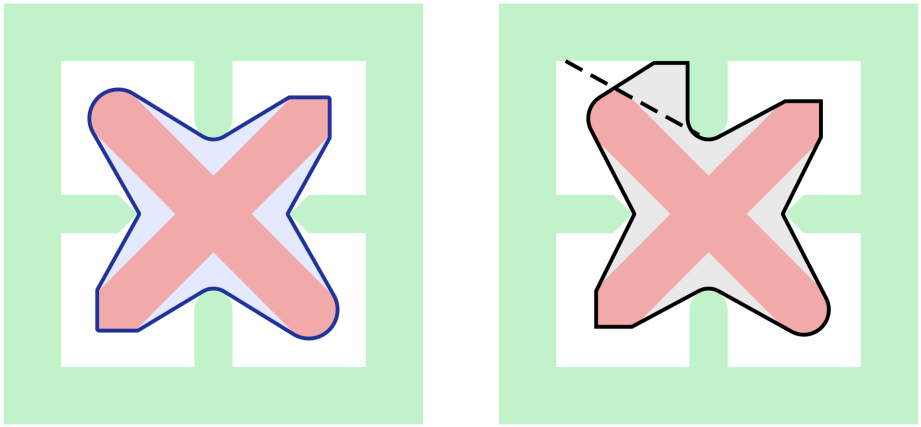


FIGURE 7.1. The dark blue hull boundary on the left intersects the boundaries of R and G, shaded red and green, only where they are convex. The boundary of the set shaded gray on the right, which intersects the boundary of G where it is concave, can be shortened by shaving the set at the dashed line. 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.

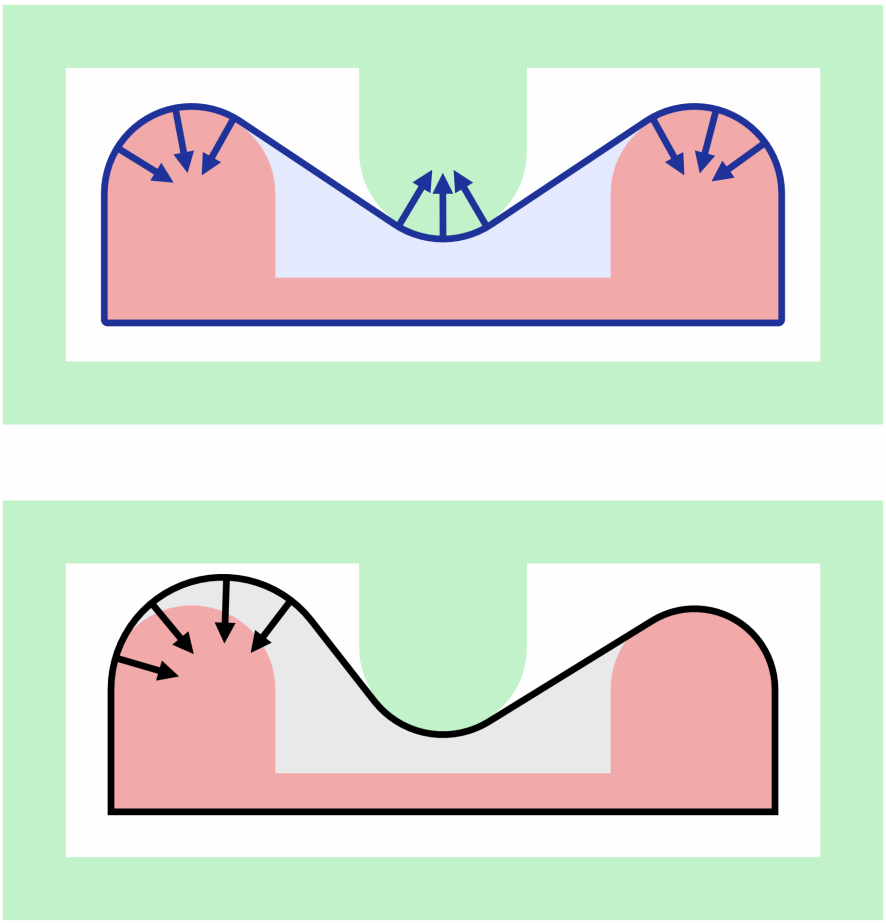


FIGURE 7.2. Above we show a hull boundary with the inward-pointing normals at several points scaled by their corresponding signed curvatures. Each scaled normal points toward the interior of either R or G, so the hull's boundary is invariant under curvature flow. At the bottom, we show a set in gray 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.

a two-dimensional tight hull is then invariant under constrained curvature flow, so it locally minimizes length [**32, 75**].

Because it is invariant under curvature flow, the unsupported subset of the boundary of a two-dimensional hull has zero curvature and 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 7.3.) This implies that the maximum absolute value of

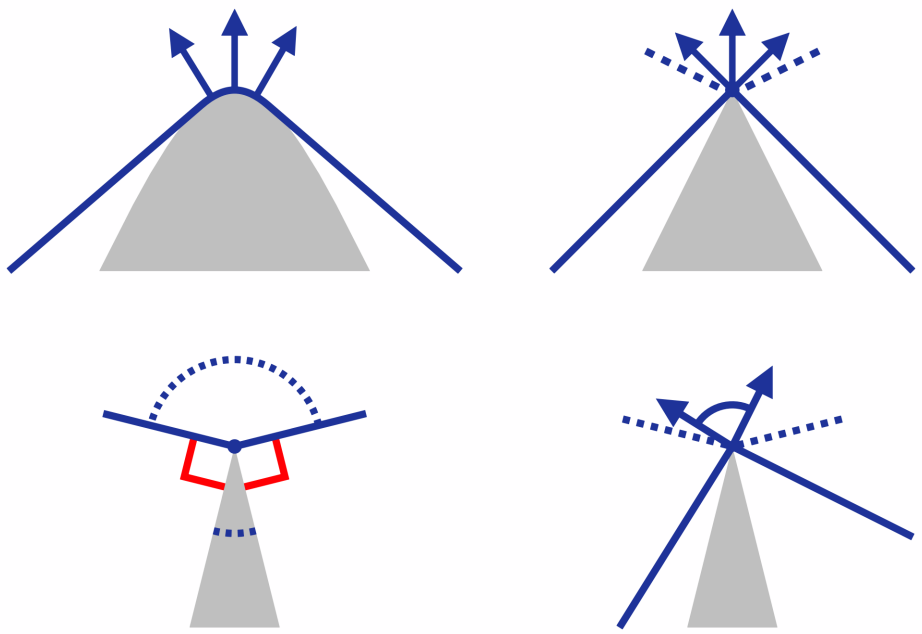


FIGURE 7.3. On the upper left, the hull boundary is smooth, and its normal variation precisely equals the normal variation of the supporting set in gray. On the upper right, we use blue arrows to indicate the fan of normals to the hull boundary at the marked vertex, and we use two dashed blue segments 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.

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

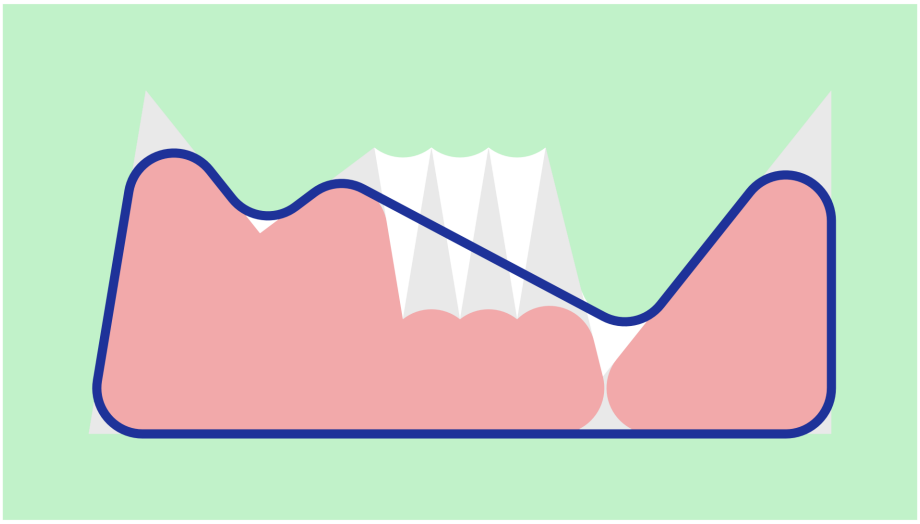


FIGURE 7.4. We show the opening of a polygon in red and the opening of its complement in green, with the difference between the polygon and its opening shaded gray. We outline the boundary of the polygon's tight blend in blue. The blue curve bounds the tight hull of the red set relative to the green set.

Definition 7.1. A *tight r-blend* of S is a tight hull of $S\circ_r$ relative to $S^c\circ_r$ (Figure 7.4.)

The boundary of a tight blend lies in the mortar, which is the set difference $S\bullet_r \setminus S\circ_r$ between the closing and opening of S.

Because 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 7.5.) Because a tight blend is a tight hull, its normal variation is no greater 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 to $\frac{1}{r}$, so the magnitude of the curvature along the boundary of a tight r-blend is also less than or equal to $\frac{1}{r}$.

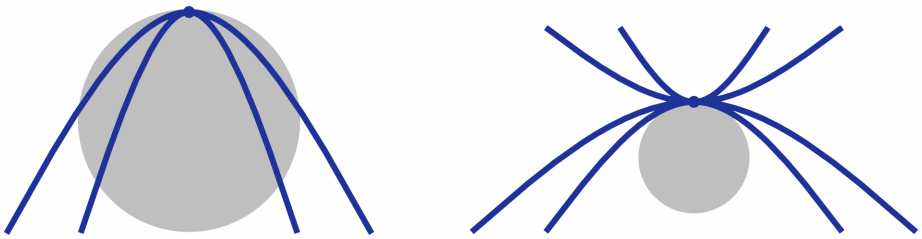


FIGURE 7.5. On the left, we show two curves in blue and a gray disk of radius r passing through the same point with the same normal. Both of the curves have a curvature greater than $\frac{1}{r}$ at the point indicated by the blue dot, and so they locally lie within the disk. On the right, we show a 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 curves is negative,) so none of them intersect the interior of the gray disk in a small neighborhood of the point.

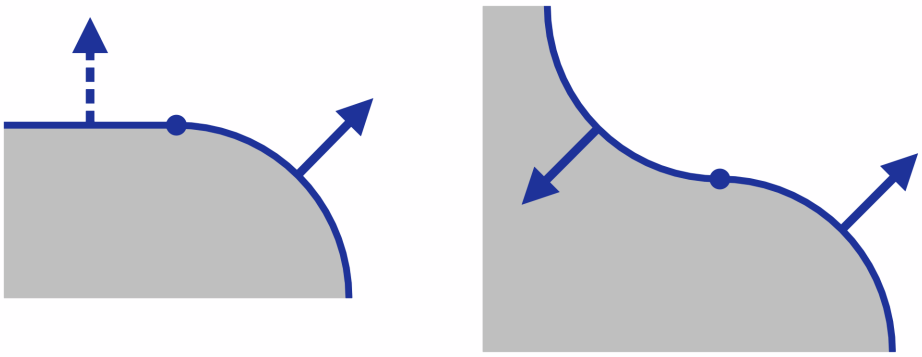


FIGURE 7.6. On the left, we show a line segment tangent to an arc from a circle of radius r at a point indicated by a blue dot on the boundary of a set shown in gray. Although the normals to the set's boundary are continuous, they are not differentiable. The normals' rate of change is zero over the line 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.

Under the definition of curvature as the rate of change of the normal as a function of arc length, the curvature is only defined where the normal field is differentiable. A tight blend may contain points where the normals are continuous but not differentiable (Figure 7.6,) so we can use an alternative definition of curvature that is applicable

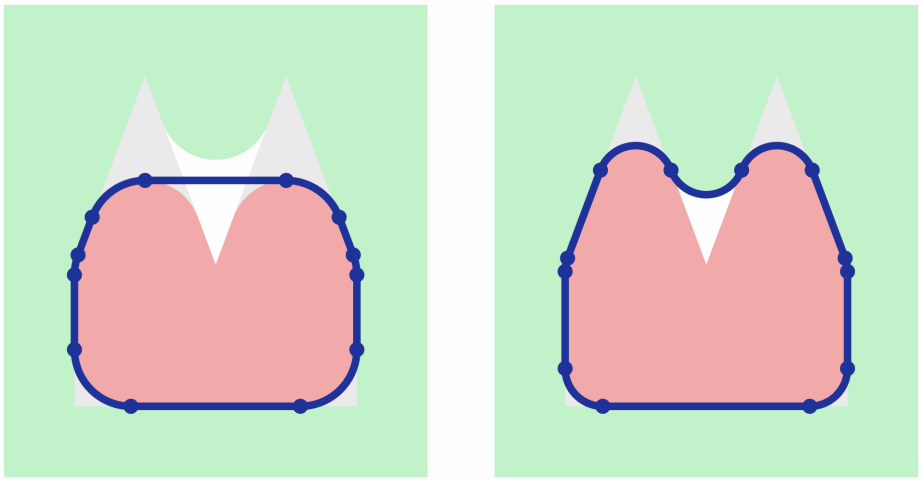


FIGURE 7.7. A tight r-blend of a polygon typically consists of an alternating sequence of line segments and arcs from circles of radius r, as shown on the left. Each segment is tangent to two arcs, one at each endpoint. On the right we show a singular situation where two arcs meet with normal continuity. We can consider there to be an edge of length zero connecting the arcs.

wherever the normal field is continuous. With this definition, we can assert that every point on the boundary of a tight r-blend has a curvature in $\left[-\frac{1}{r}, \frac{1}{r}\right]$.

7.2. Continuous tight blends of polygons

We consider tight blends of polygonal data to be of special interest. In part, this is because polygonal shape representations are ubiquitous, and in part it is because a tight blend 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 7.7.)

For every polygon with a manifold boundary, there is a value of r small enough so that the polygon's r-blend 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, the arc uniformly distributes the change in the normal over a curve whose length is proportional to r (Figure 7.8.)

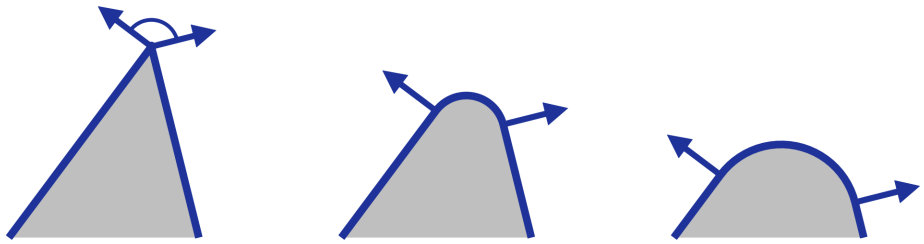


FIGURE 7.8. On the left, we show the change in normal direction that occurs at a vertex as we move from one edge incident on the vertex to the other. In the middle, we replace the vertex with a small arc, distributing the change in direction over the arc. On the right, we replace the vertex with a larger 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.

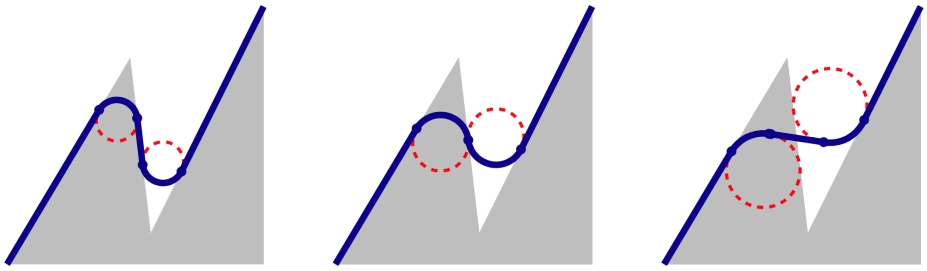


FIGURE 7.9. On the left, we show a portion of a gray polygon's tight blend with a small radius. We draw the boundary of the tight blend in blue, with dots marking the endpoints of two arcs and the edge connecting them. In the middle, we increase the radius used for a blend. This makes the arcs longer, and their endpoints slide toward each other along the polygon's edge, eliminating the edge that connected them on the left. On the right, we show how further increasing the tight blend's radius introduces an edge into the blend's boundary that does not lie on an edge from the input.

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 the edges that connect the arcs decrease (Figure 7.9, left.) At some value of r, two arcs meet and an edge disappears (Figure 7.9, center.) As r continues to increase, a new edge appears (Figure 7.9, 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.

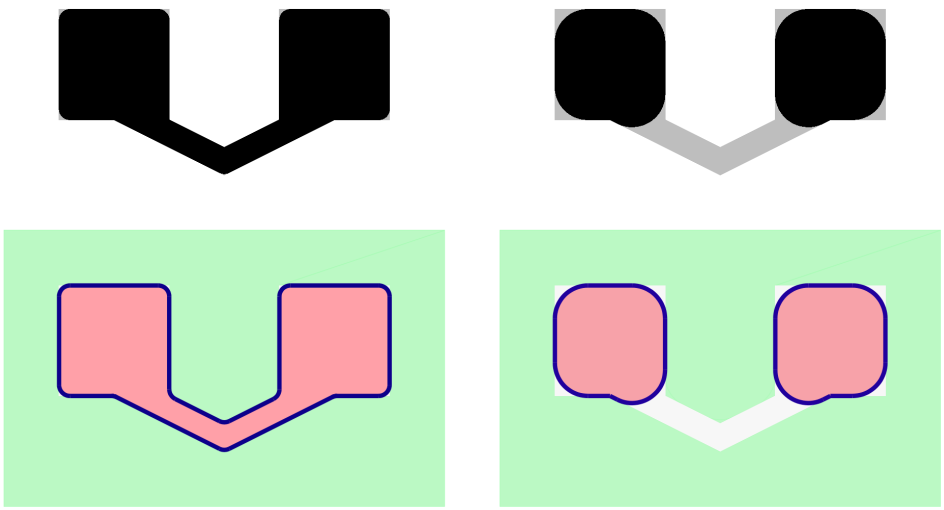


FIGURE 7.10. On the top left, we show a polygon in gray with its morphological opening by a small disk superimposed in black. On the top right, we show the same polygon with a larger-radius opening superimposed. On the bottom left, we show the boundary of the tight blend of the polygon with a small radius 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 a tight blend with a larger radius. Although the tight blend on the bottom left has one connected component, the tight blend on the bottom right has two.

Morphologically opening a set may change its topology. Accordingly, there may be a value of r where the topology of a set's r-blend changes, producing a discontinuous change in its boundary (Figure 7.10) Over intervals of r free of topological changes, blending defines an isotopy. For a large class of sets, the number of topological changes is finite, so that we can characterize blending as piecewise continuous.

7.3. Tight blends simplify normal fields

Over an interval of values of r where the boundary of a tight r-blend 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.

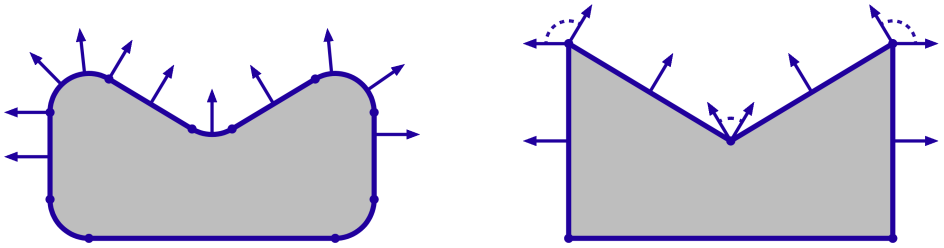


FIGURE 7.11. The normal field for the set produced by a tight blend on the left is defined by the normals to its bounding arcs and edges. Because the edges of the polygon on the right meet at singular points, each polygon vertex possesses a set of normals covering the arc swept when we rotate the normal to one of its incident edges until it aligns with the normal to the other edge.

Over an interval where blending 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.

To define the Gauss offset of a set $S \subseteq \mathbb{R}^d$, we need S to have a well-defined normal field. A tight blend's normal field is given by the outward-pointing normals of its bounding arcs and edges. Because a tight blend 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 7.11). We can 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 7.12.)

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

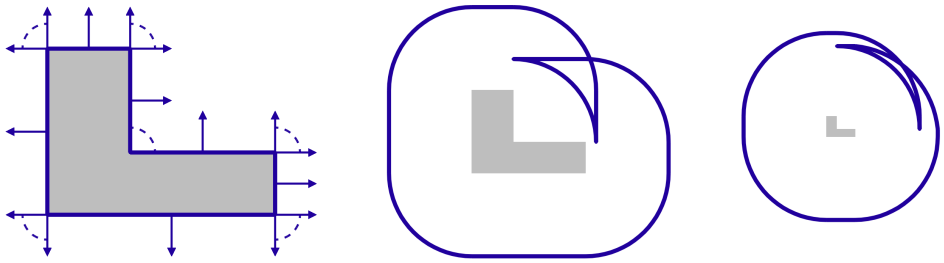


FIGURE 7.12. On the left, we show a nonconvex polygon in gray with its normal field indicated in blue. In the middle, we show the polygon's $\frac{1}{2}$-Gauss offset, which self-intersects due to the polygon's concave vertex. On the right, we show the polygon's ϵ-Gauss offset as ϵ goes to zero. The 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.

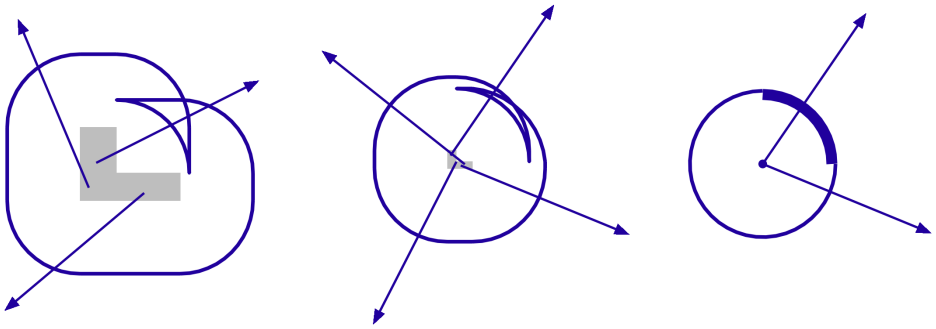


FIGURE 7.13. If we pick a ray with its tail at a point in a set and count the number of times the ray intersects the set's ϵ-Gauss offset, the result depends on the tail position and on the precise value of ϵ when ϵ is large. As ϵ decreases, the set of tail positions converge to a point and the geometry 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.

components from the set's boundary whose normals have that direction (Figure 7.13) We refer to this function as the polygon's normal count.

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

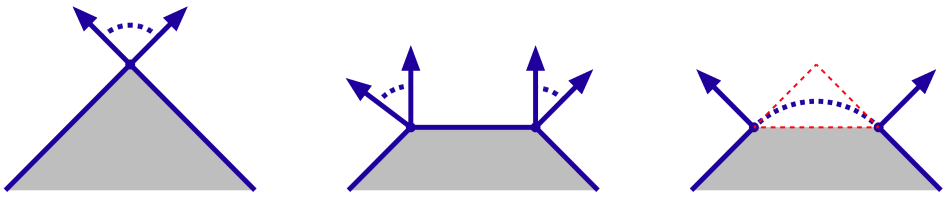


FIGURE 7.14. On the left, we show a corner of a polygon. Slack is concentrated at the corner and is equal to the angle between the two edges incident on the vertex. In the middle, we cut the corner, which divides the slack into two equal parts located at the shown vertices. We are given two edges with a gap between them on the right. Any convex function defined over the dashed red line that remains within the dashed 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.

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 blending as long as blending 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 7.8. Every boundary segment connecting two points has the same slack, provided the change in normal direction over the segment is monotonic (Figure 7.14.) 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 7.15.) 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 combine blending's slack-reducing and curvature-limiting properties to conclude that a tight blend of 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.

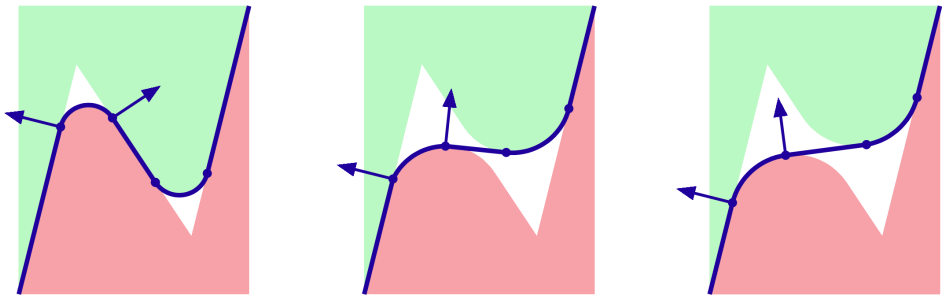


FIGURE 7.15. On the left, we show a portion of a tight r-blend of a polygon. We mark the endpoints of two arcs from the blend and the tangent edge connecting them with blue dots. Two arrows indicate the change in normal direction over one of the arcs. In the middle, we show a tight blend of the same 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 tight blend. 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 8

MEDIAL COVER

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 **[70]** 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 (**[16, 47, 70]**, Figure 8.1.)

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 **[70]**, but the fact that the geodesic hull is connected makes it 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 (Figure 8.2.) The number of tight covers may be exponential in the number of components in R (Figure 8.3,) 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 8.4, **[36]**.)

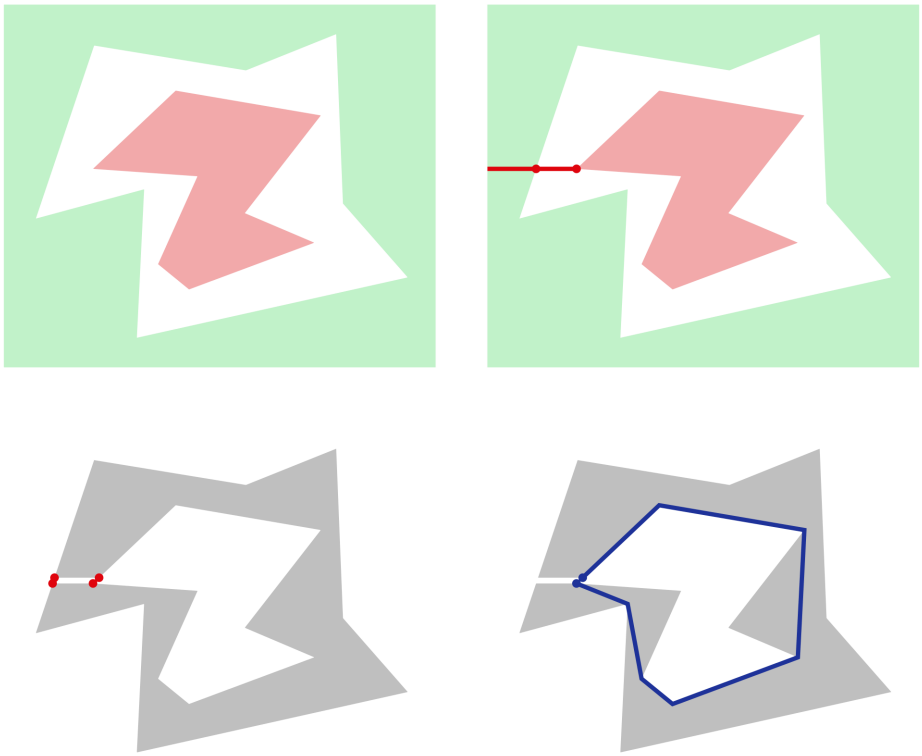


FIGURE 8.1. On the upper left, we show a set R in red and G in green, where R and G^c are simple polygons. The set $(R \cup G)^c$, which is the gap between R and G, is isotopic to an annulus. On the upper right, we cut the annulus horizontally with a ray extending leftward from the leftmost vertex v 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 $\overrightarrow{vp}$ 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, we use the medial axis to specify a set of points that we add to R and G (Figure 8.5.) 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 points. It can also

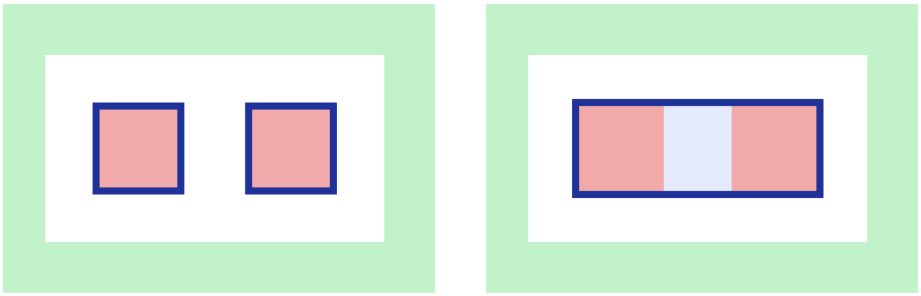


FIGURE 8.2. On the left, we outline in blue the boundary of a tight cover of a set R consisting of two squares relative to a set G equal to the complement of a rectangle. The cover is equal to R and has two connected components. On the right, we outline the boundary of a tight cover of R relative to 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.

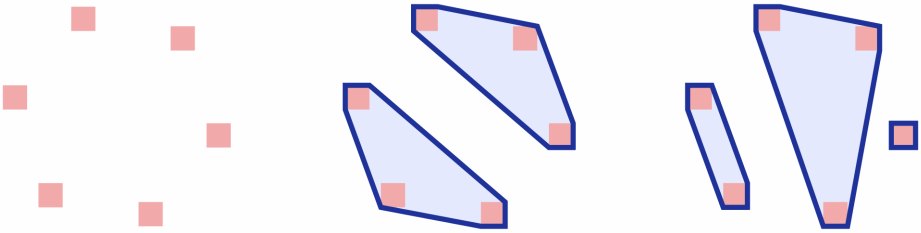


FIGURE 8.3. On the left, we show several small connected components of R arranged at points on a circle. Taking G to be empty, we show two tight covers of R relative to G at the middle and right. The number of covers is not bounded by a polynomial function of components of R.

be characterized as the tight cover whose boundary is isotopic to the set of points equidistant from R and G (Figure 8.6,) 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; 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, implying that the complement of the convex hull of G relative to R is the smallest (Figure 8.7.)

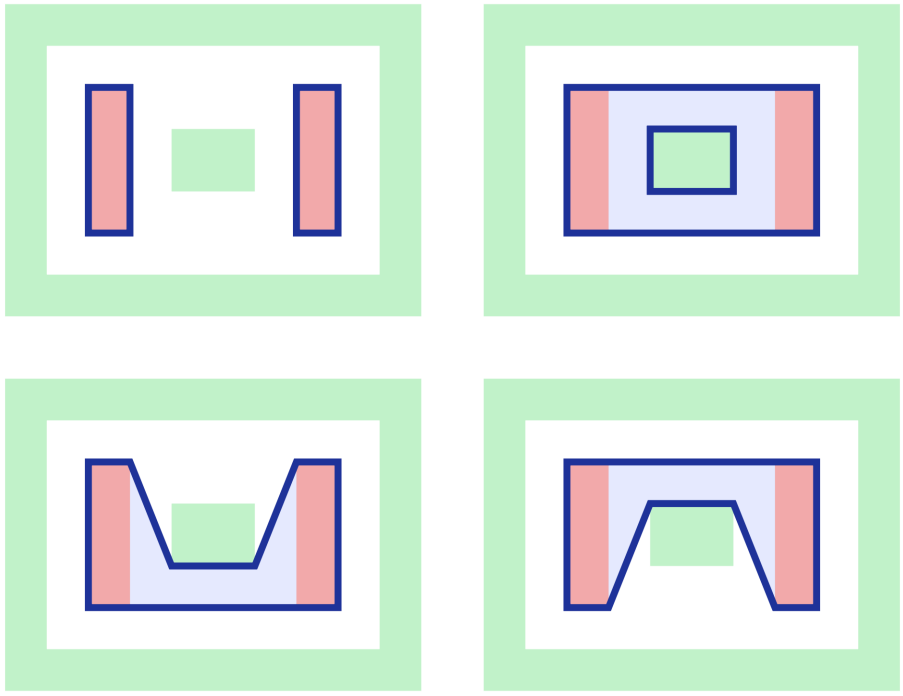


FIGURE 8.4. We show four tight covers of R relative to G. The cover in the upper left consists of two connected components. The other three covers each consist of a single component, although the component from the cover at the upper right has a hole. The two covers at the bottom are isotopic to each other in the plane, but they cannot be deformed into one another if we constrain the deforming set to contain R and exclude G.

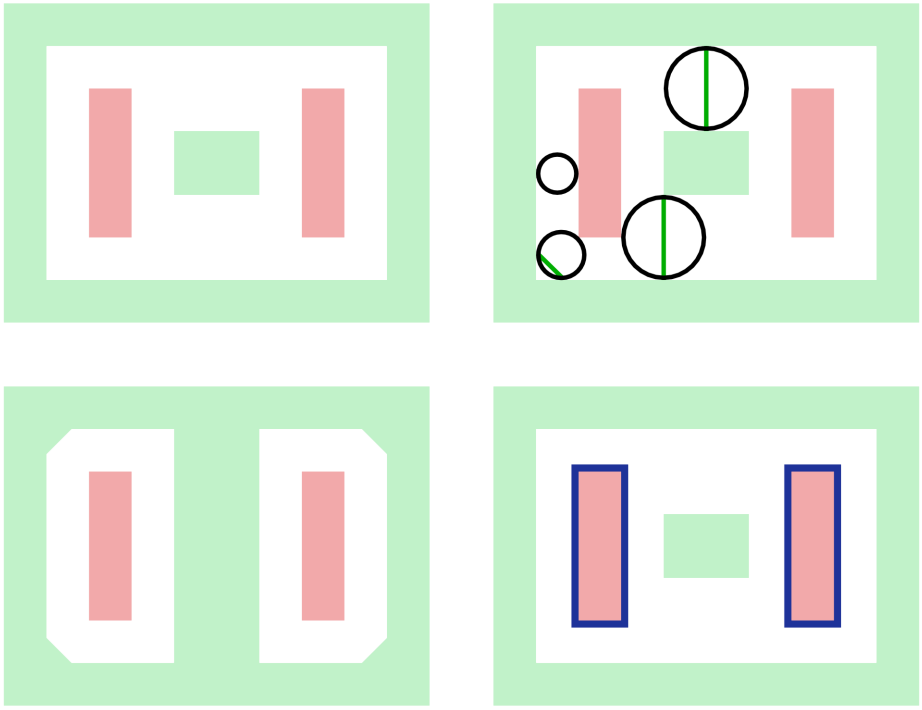


FIGURE 8.5. In the upper left, we show the sets R and G from Figure 8.4 in red and green, respectively. In the upper right, we show four of the maximal disks centered on the medial axis of $(R \cup G)^c$. For each disk, we take the convex hull of the points where it touches R and add the result to R. 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 tight cover of R relative to G, which we obtain by constructing the shortest loop in each annulus. This medial tight cover consists of two components. It is equal to the tight cover shown in the upper right of Figure 8.4, and it is different from the tight hull of R relative to G.

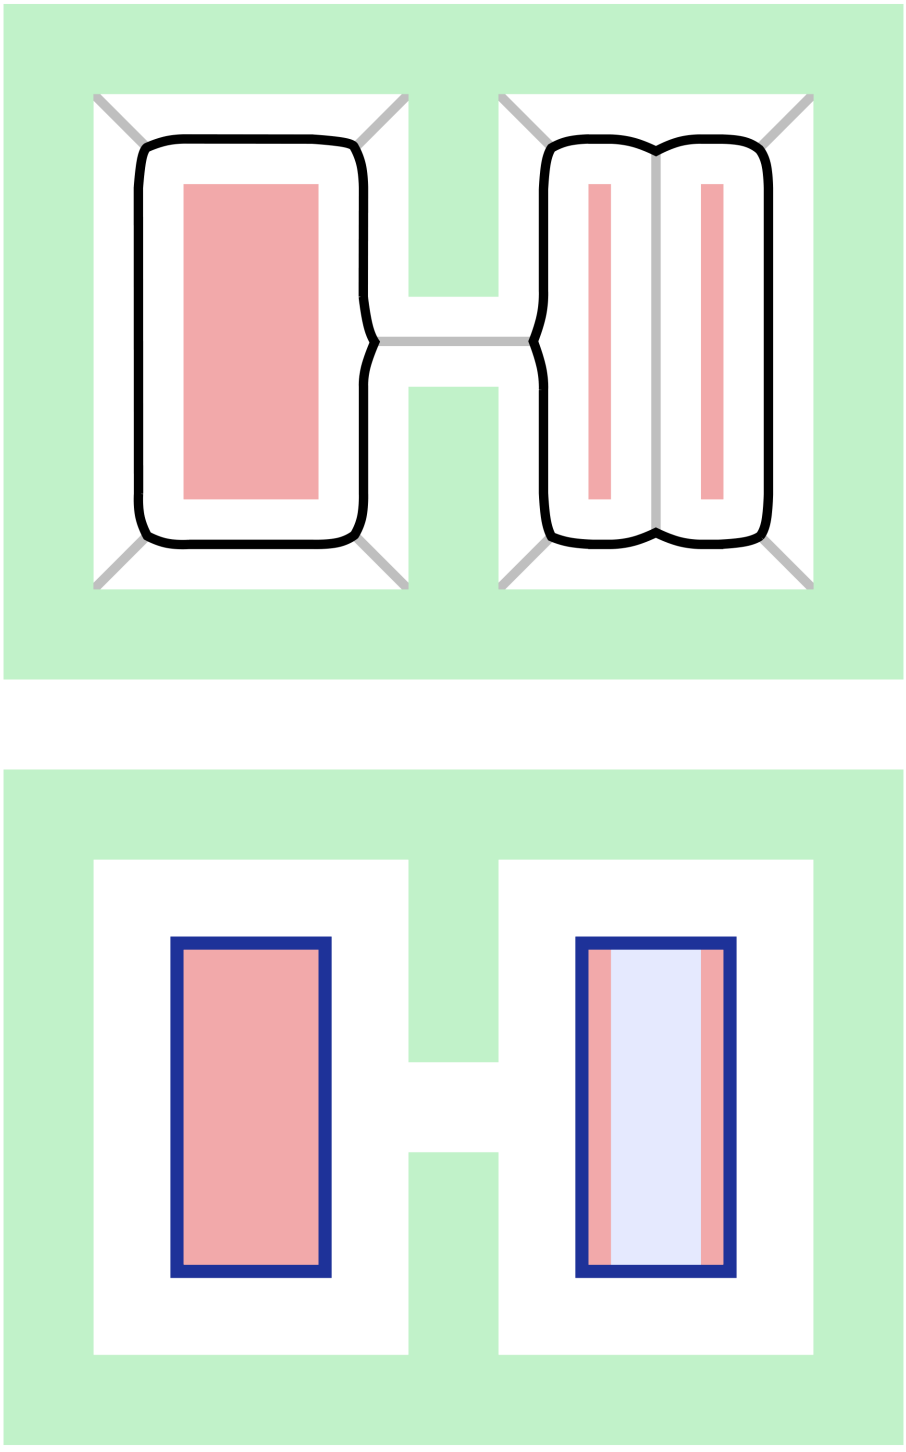


FIGURE 8.6. Above, we show R and G with the medial axis of $(R \cup G)^c$ in gray and the points equidistant from R and G (which form a subset of the the medial axis) in black. Below, we show the medial cover of R relative to G in blue. The black loops above may be deformed to the blue loops on the right without intersecting the interiors of R or G.

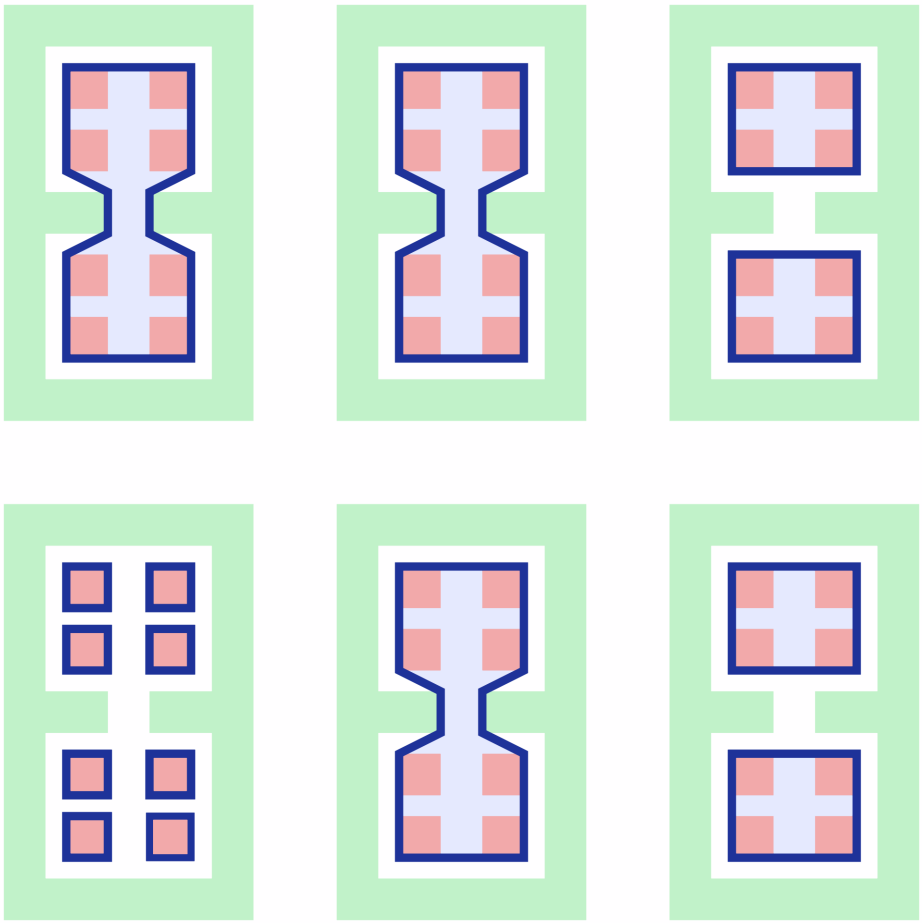


FIGURE 8.7. At the top left, we show the convex hull of R relative to G, which consists of a single large component. At the top middle, we show the tight hull of R relative to G, which in this case is identical to the convex hull of R relative to G. At the top right, we show the medial cover of R relative 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 9

IMPLEMENTATION

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 9.1.) 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

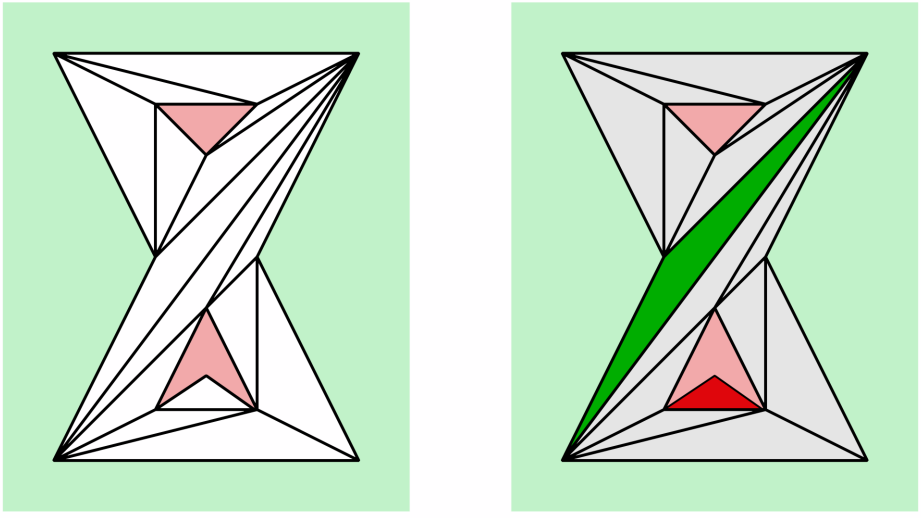


FIGURE 9.1. On the left, we show a triangulation of the space between R and G. On the right, we shade gray the triangles whose vertices lie on the boundaries of both R and G. We shade the triangle whose vertices lie on R red, and we shade the triangle whose vertices lie on G green.

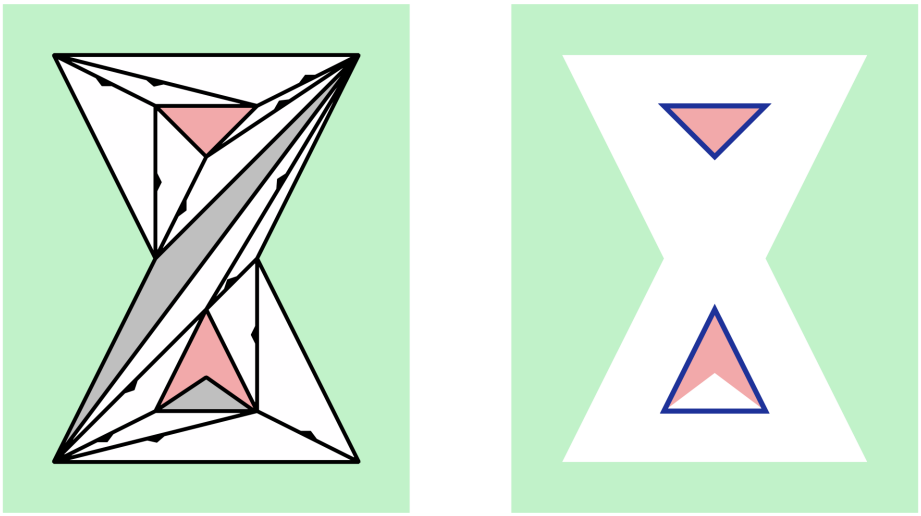


FIGURE 9.2. On the left, we illustrate how each mixed triangle connects to two mixed triangles, one across each of its two mixed edges. We shade the solid triangles gray. Using a consistent triangle orientation, we can assign each mixed triangle a predecessor and successor. We indicate this with a black 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.

of each annulus relative to the unbounded component can be computed in linear time, as we explained in Chapter 8. The annuli decompose the medial cover problem into the problem of computing a linear number of 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 9.2.) 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.

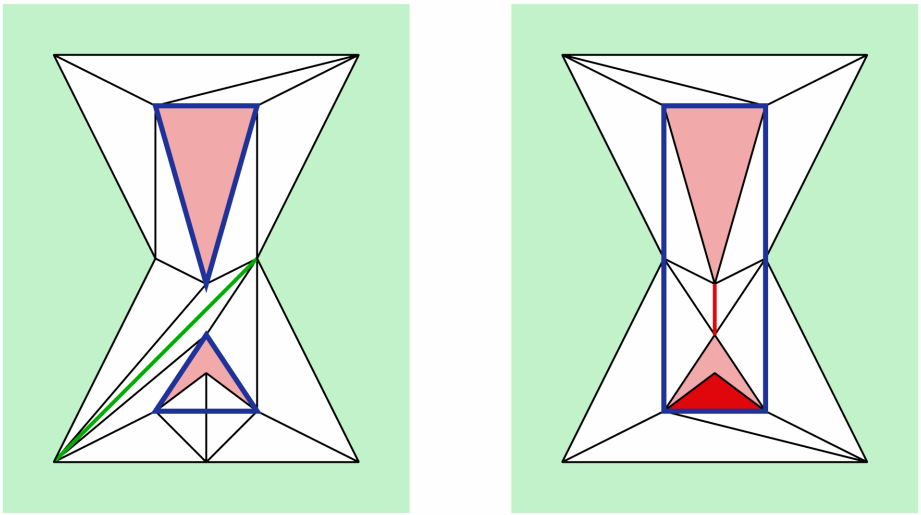


FIGURE 9.3. On the left, we show a tight cover of R relative to G constructed from an arbitrary triangulation. We shade the solid edge connecting two vertices of G green and outline the hull boundary in blue. On the right, we show the constrained Delaunay cover of R relative to G, which we construct 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.

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 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 a tight blend with a progressively increasing radius.

We can address this limitation by using a geometrically-based triangulation, such as a constrained Delaunay triangulation [**19**]. 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 9.3.) 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

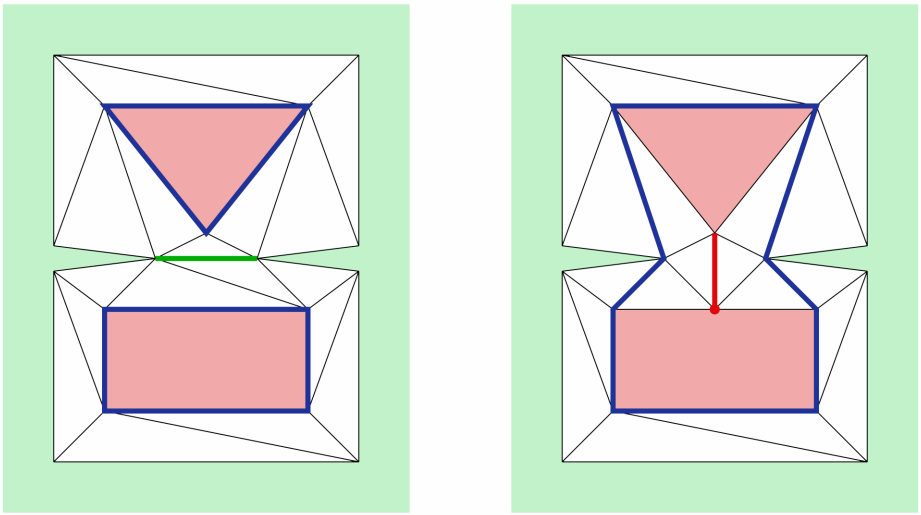


FIGURE 9.4. On the left, we show a constrained Delaunay cover constructed from vertices located at the corners of R and G. On the right, we introduce a vertex into one of the edges bounding R. The constrained Delaunay triangulation using that vertex is different from the one that does not, yielding a different constrained Delaunay cover.

to G may change although R and G do not (Figure 9.4.) 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.

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 of 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 9.5.)

The circumcenters of the large triangles approximate bifurcation points of the medial axis (Figure 9.6.) 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 9.7.) The constrained Delaunay cover of this modified input is the same as its medial cover; inserting additional vertices

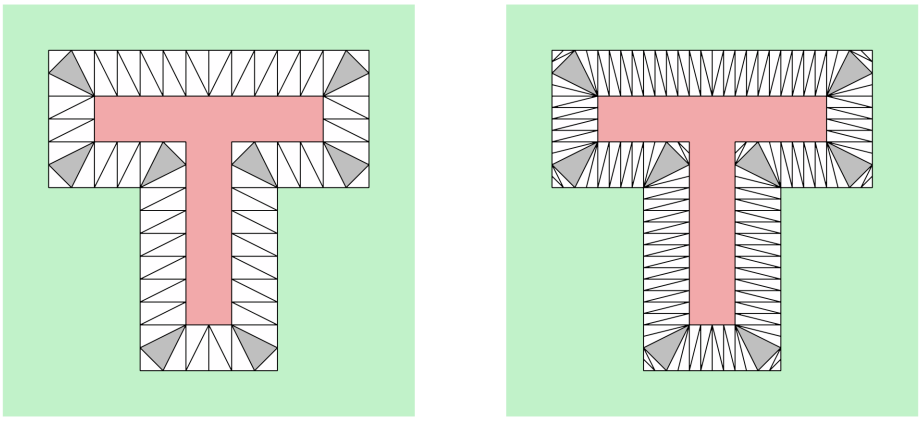


FIGURE 9.5. On the left, we show a constrained Delaunay triangulation of the space between uniformly sampled polygonal sets R and G. We shade gray the triangles that only touch the boundaries of R and G at their vertices. The gray triangles are located at corners of R and G, and they are slightly larger 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.

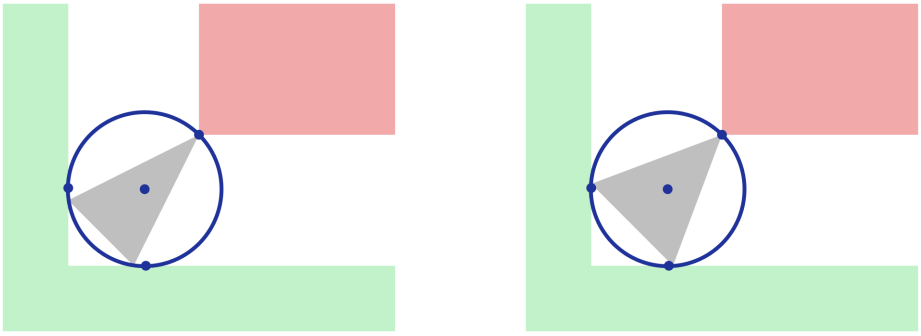


FIGURE 9.6. On the left, we zoom in on a portion of the sets R and G shown in Figure 9.5. The vertices of the large gray triangle are close to the points where a maximal disk centered on a bifurcation point of the medial axis touches the polygon. We outline the maximal disk in blue and mark its center 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.

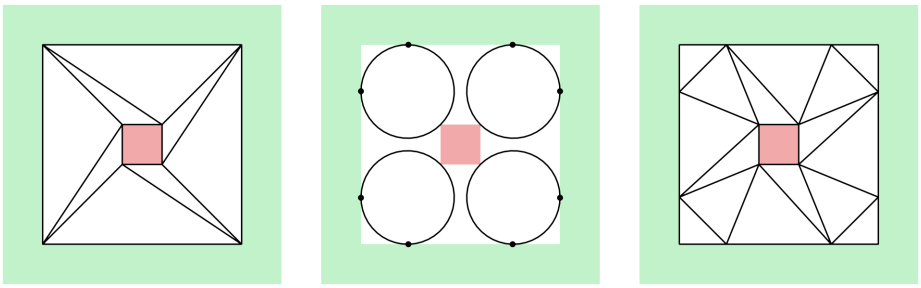


FIGURE 9.7. On the left, we show set R and G in red and green, respectively, with a constrained Delaunay triangulation of the space between R and G in black. In the middle, we outline the maximal disks centered on the bifurcation points of the medial axis of $(R \cup G)^c$ in black, and we indicate the 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.

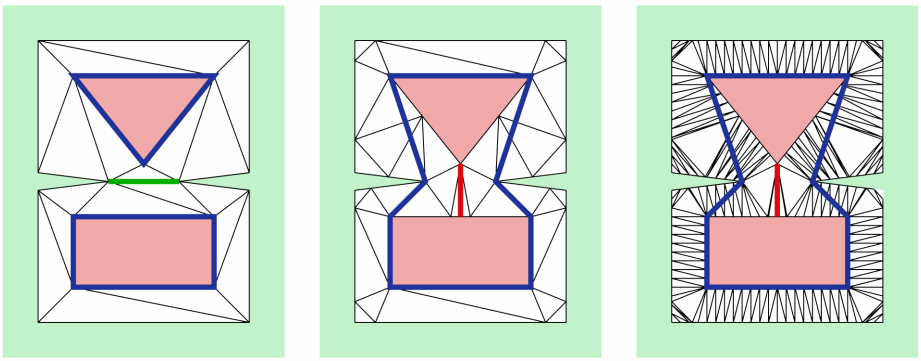


FIGURE 9.8. On the left, we show a constrained Delaunay triangulation of vertices located at the corners of R and G. We show the cover obtained from this triangulation in blue. In the middle, we add additional vertices at the contact points of the maximal disks centered on the bifurcation points of the 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.

has no effect (Figure 9.8.) 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 9.9.)

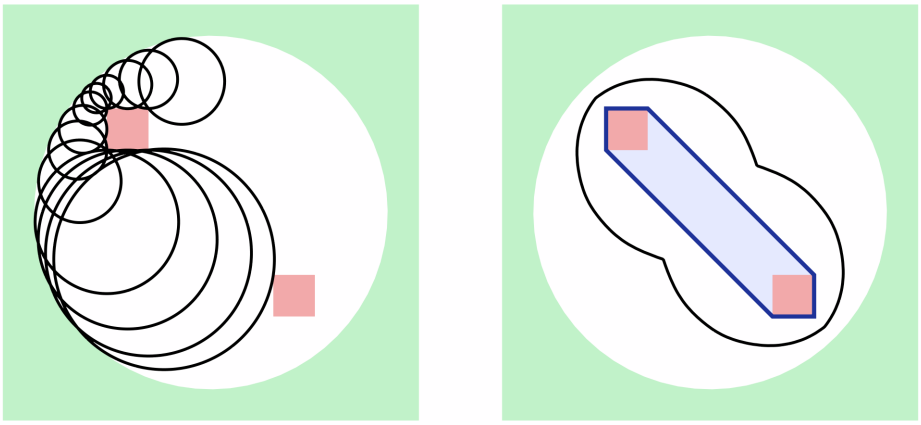


FIGURE 9.9. If we have an algorithm for computing the medial cover, we can compute the convex hull of a nonempty set R by taking G to be the complement of any closed disk containing R. Then no maximal disk in the gap between R and G touches G in more than one point. We illustrate this on the left, shading 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.

Rather than compute a tight r-blend, we compute an approximate medial blend. The medial r-blend 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 blends with disks, 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 9.10,) 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 9.11.) Although an approximate medial blend does not have bounded curvature, it does have bounds on the maximum angle change at a vertex and over a boundary interval (Figure 9.12.)



FIGURE 9.10. 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.

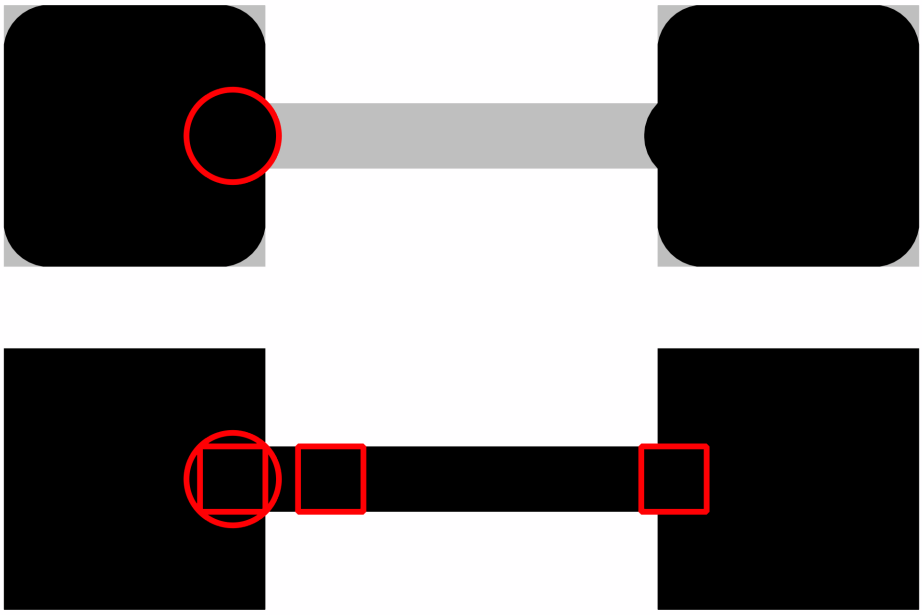


FIGURE 9.11. Above, we show a set S in gray with the opening of S using a disk of radius r overlaid in black. We outline the disk in red. Below, we show the opening of S using a square inscribed in the disk. Unlike the opening of S with the disk, the opening of S with the inscribed square is identical 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.

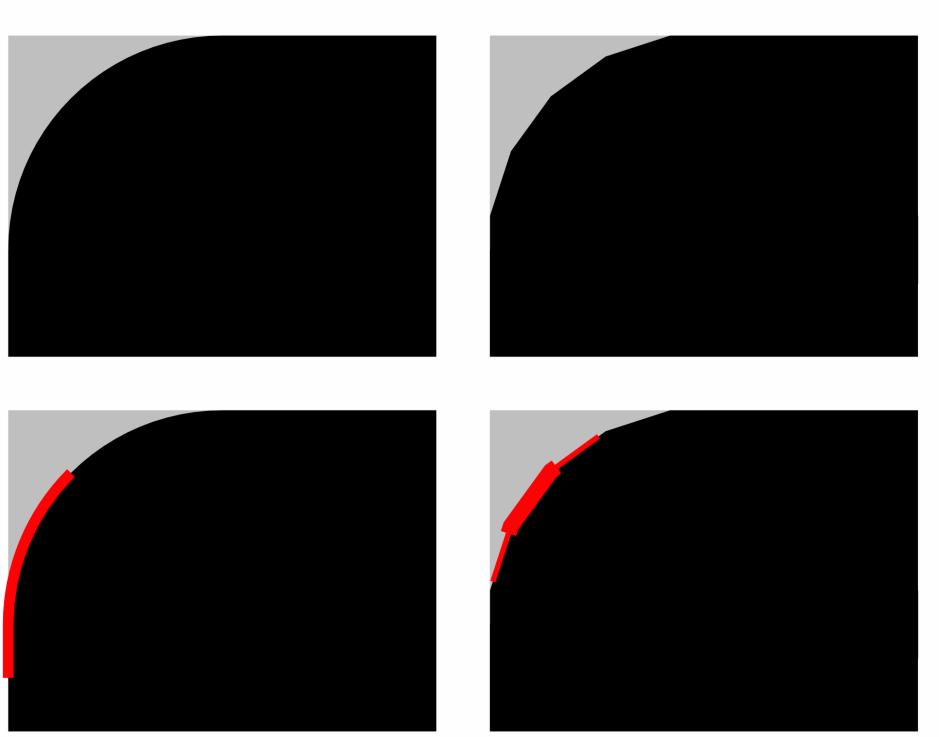


FIGURE 9.12. In the upper left, we illustrate how a tight r-blend replaces a corner of a square with a circular arc of radius r. In the upper right, we show how an approximate blend replaces the corner with a piece of a regular polygon. A total angle change of $\frac{\pi}{2}$ occurs between the endpoints 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 arc length 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 an approximate blend 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 an approximate r-blend is no greater than $\frac{1}{r}$ times the interval's length plus an error inversely proportional to n.

CHAPTER 10

APPLICATIONS

10.1. 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 [**65**]. 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 10.1.) 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 every square bounded

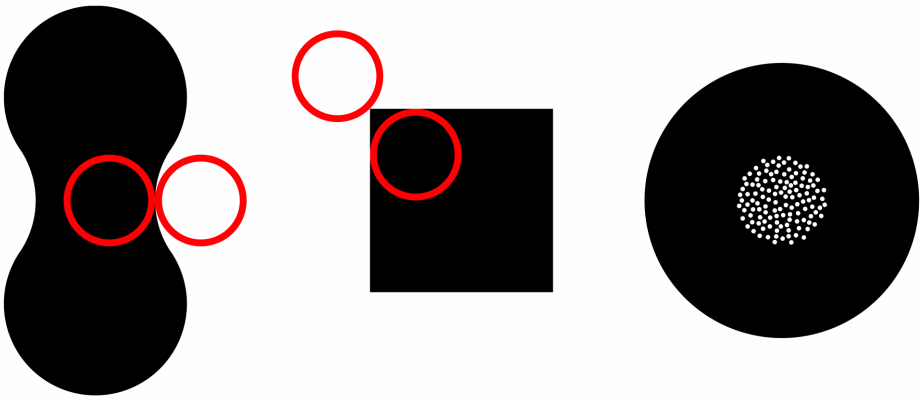


FIGURE 10.1. On the left, we show a morphologically r-regular set. It has a manifold boundary and its interior and complement can be expressed as unions of balls of radius r. The interior of the square in the middle cannot be expressed as a union of balls of radius r, so it is not r-regular. The set 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.

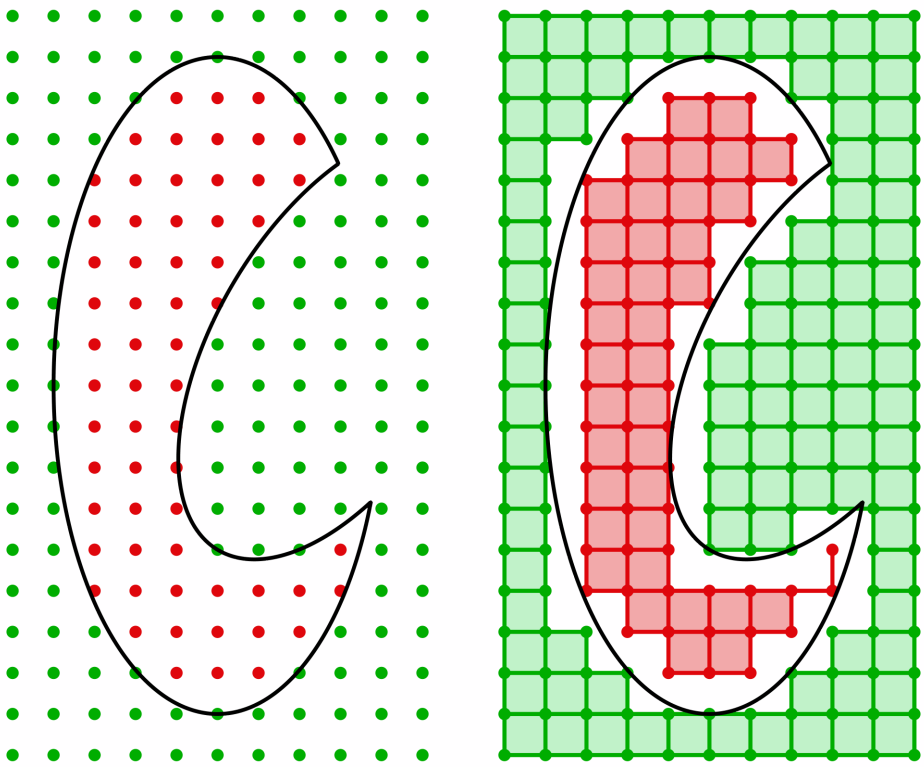


FIGURE 10.2. On the left, we outline a set S in black. We show a rectangular grid of samples, with samples inside of a set S colored red and samples outside of S colored green. On the right, we shade red the edges and squares with red vertices, while we shade green the edges and squares with green vertices.

by four red edges. We likewise color green the edges connecting green vertices and the squares bounded by green edges (Figure 10.2.)

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 10.3; [**65**].) 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 [**39, 52**] obtained convergent estimates of the d-dimensional measure of a d-dimensional set in the 19th century (Figure 10.4,) convergent $(d-1)$-

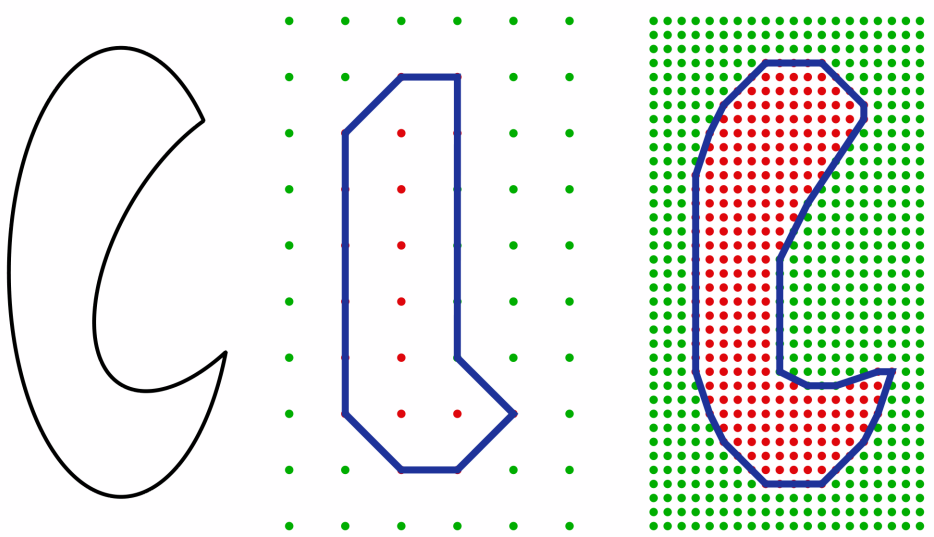


FIGURE 10.3. On the left, we outline a set S with boundary length of 91.0. In the middle, we show the relative convex hull's reconstruction of the boundary of S from a coarse sampling in blue. The length of the reconstructed boundary is 70.6. On the right, we show the relative convex hull's boundary 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 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 10.5.)

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. 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 fundamental theorems

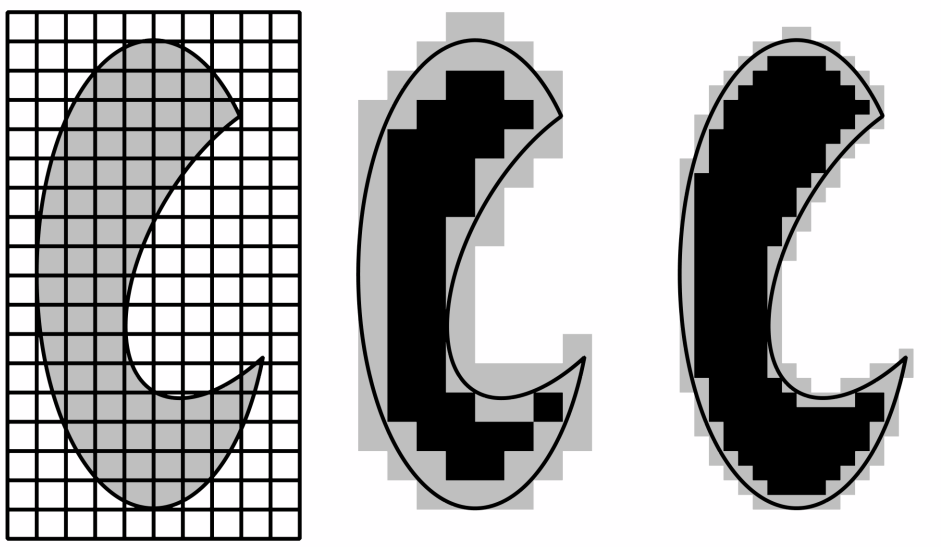


FIGURE 10.4. Peano and Jordan define the area of a planar set S by overlaying a square grid on it, as illustrated on the left. Taking the area of a square to be the product of its edge lengths, they define the outer measure of S as the sum of the areas of the squares intersected by S, while the inner 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.

of geometry, such as the proposition that a circle is the set of minimum perimeter out of all sets with a given area [**43**].

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 [**43**]. For example, the surface area of an ellipsoid or a Boolean

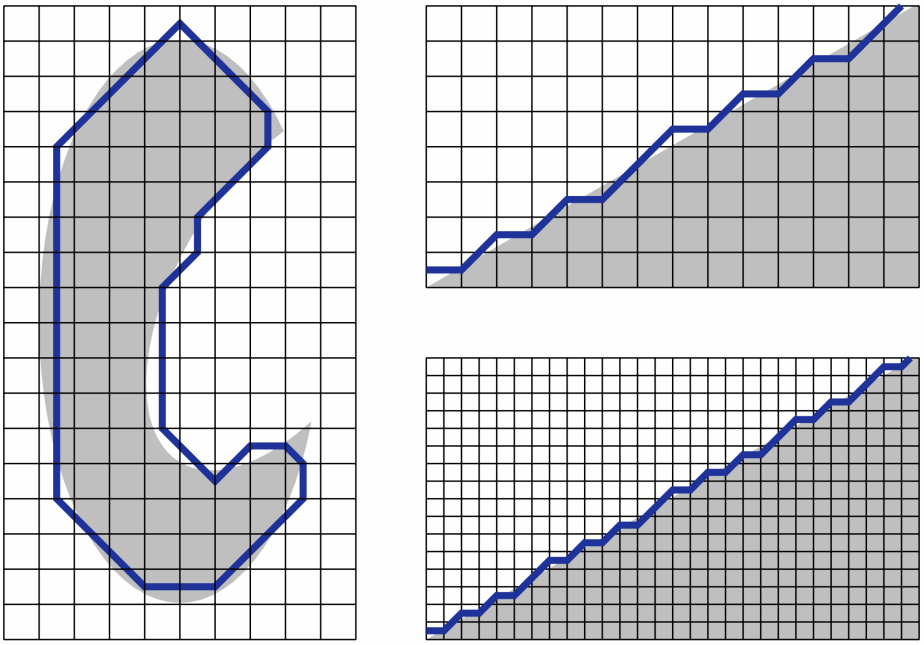


FIGURE 10.5. On the left, we show a midpoint reconstruction of the boundary of the set S from Figure 10.3 in blue. Compared to the boundary reconstruction obtained from the relative convex hull, the midpoint reconstruction is more jagged, which increases its length. To highlight this behavior, we consider 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.

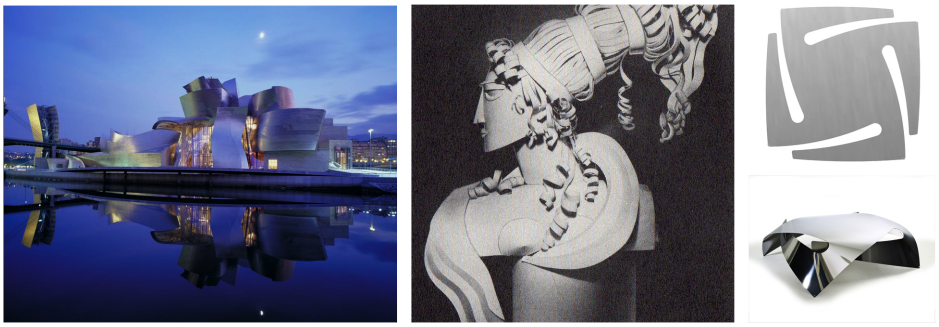


FIGURE 10.6. On the left, we show an image of the facade of the Guggenheim at Bilbao [**61**]. In the middle, we show a paper sculpture used in a perfume store display [**31**], and on the right we show a coffee table folded from a laser-cut sheet of stainless steel [**40**].

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 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.

10.2. 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 10.6; [**30, 41, 54, 55, 57**].) While several papers address the problem of computing a developable patch interpolating a given curve (Figure 10.7; [**30, 57**],) tight hulls may be the first work on developable surfaces

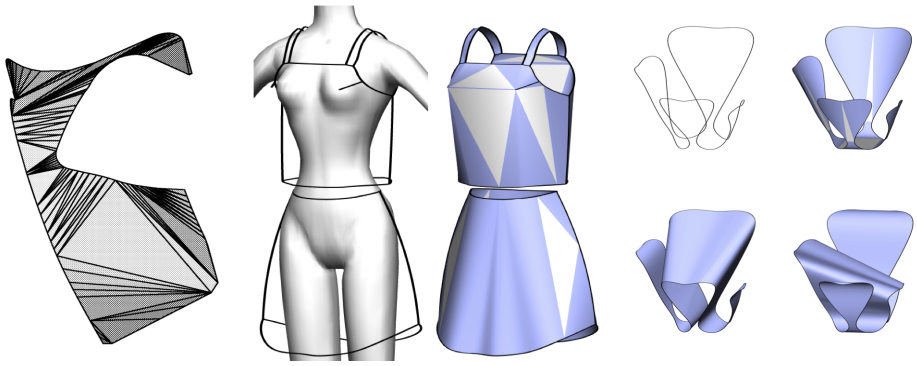


FIGURE 10.7. On the left, we show a triangulation of a boundary curve [**30**] that estimates the geometry assumed by a sheet of metal pressed by a machine. The center image [**57**] illustrates garment design by interpolating seams with developable patches. The right images [**57**] show how multiple patches can interpolate the same boundary curve, providing us with motivation to determine which patch is "best."

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 include and exclude. For instance, we could produce an executable illustration of the thought experiment described in Subsection 6.1, 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 [**51**], but with linear gradients along the rulings of the developable surfaces we construct.

10.3. Blending

We anticipate two significant applications of tight blends. The first is to solid blending. In two dimensions, applying a tight blend to a set limits the magnitude of its curvature, making it globally smooth. By contrast, it is impossible to guarantee that a set is made smooth by morphological opening or closing, or with any composition of

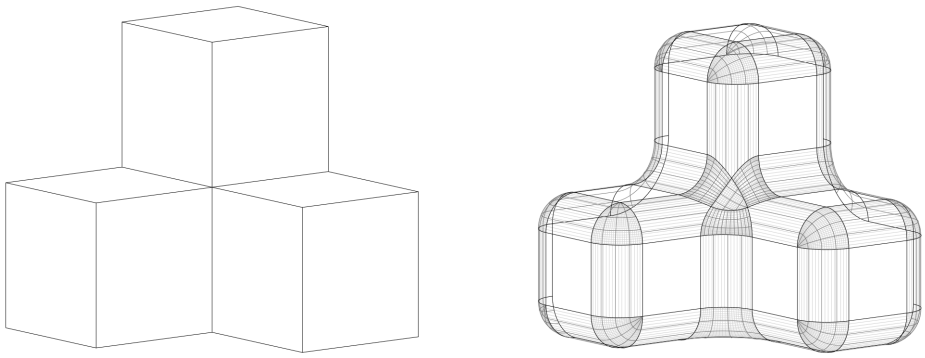


FIGURE 10.8. On the left, we show a wireframe image of a solid S with a singular saddle point. On the right, we show a schematic illustration of the tight blend of S with a small radius, which has a unique normal at every point.

the two. At present, we know little about the properties of tight blends in three and higher dimensions, but tight blends apparently guarantee the existence of a bounded sectional curvature at every boundary point. Even in the absence of a guaranteed curvature bound, tight blends may produce useful blends in three dimensions, just as opening and closing produce useful blends despite the fact that they do not guarantee smoothness. Compared to a composition of opening and closing, tight blends have the advantage that they are symmetric with respect to set complement, so they have no bias toward either enlarging or reducing a set.

10.4. Normal field simplification

The second application of tight blends is to normal field simplification. Tight blends have potential for use as a theoretical tool. Although we can define normal fields to sets with singularities, the outward-pointing normal at a singularity is not unique. A small-radius tight blend of a set, however, is a close approximation to the set with a unique normal at every point on its boundary (Figure 10.8.) As we increase the tight blend 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 10.9.)

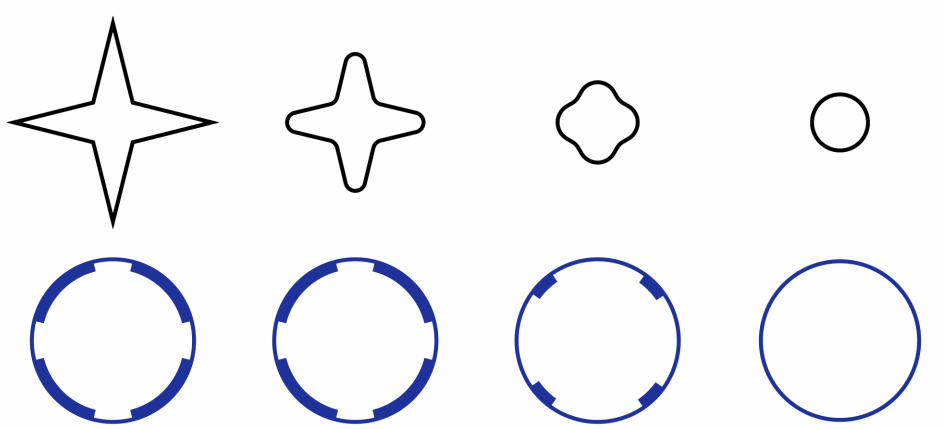


FIGURE 10.9. On the top row, we show the progressive tight blend of a four-pointed star. Beneath each star we show the star's normal field coverage. The untightened and slightly tightened stars have identical normal fields, but the normal field variation decreases as the tight blend radius further increases. Before disappearing entirely, the tight blend star is a disk, and its normal field uses every normal direction exactly once.

10.5. Further applications

Some applications of tight hulls show promise but remain speculative. An example is the use of tight hulls for the topological repair of polygonal models. In two dimensions, suppose we are given an unorganized collection of edges and vertices. We 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 10.10.) 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 tight hull: 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.

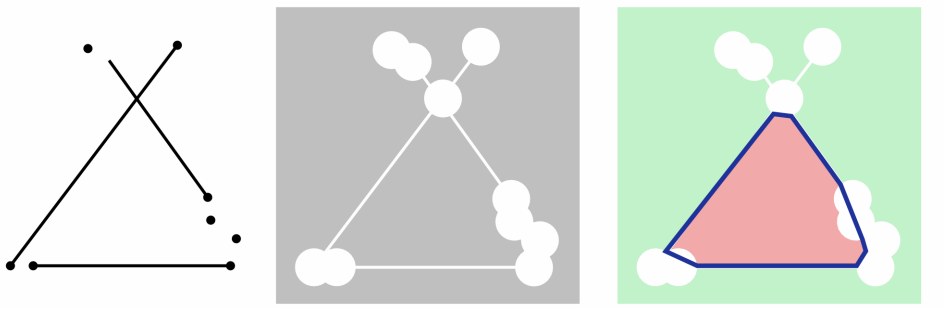


FIGURE 10.10. On the left, we show a nonmanifold soup of edges and vertices. In the middle, we shade gray the set X obtained by subtracting the union of the input with the dilation of its nonmanifold points from the plane. On the right, we designate the bounded component of X as belonging to the red 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.

CHAPTER 11

CONCLUSION

11.1. Introduction

We begin by summarizing our contributions in Chapters 3-10, including definitions, properties, implementations, and applications. Then we discuss the implications of our work, arguing the tight hull's importance in the context of mathematical morphology, fairing, and convexity. We conclude by describing two significant aspects of our future work. The first pursues efficient implementations of three-dimensional tight hulls, and the second seeks regular shapes from input shapes.

11.2. Contributions

Mason. Building on the concept of regularity, mason identifies the importance of nearly regular sets while defining irregularity, the mortar, the core, and the anticore. Mason replaces each component of the mortar by choosing a composition of opening and closing to minimize the symmetric difference between the input shape and the output. Given input represented by pixels or voxels, we use Euclidean distance mapping and traversal to apply mason in linear time.

Tightening. Tightening locally minimizes the measure of the boundary set separating the input shape's core and anticore while lying in its mortar. We construct r-tightening with mean curvature flow, guaranteeing a mean curvature bound between $-\frac{1}{r}$ and $\frac{1}{r}$. We also show that tightening allows a variety of topological choices, such as by minimizing total measure boundary or initiating mean curvature flow from the output of mason. To compute tightening, we efficiently apply level set flow to pixel or voxel representations with acceleration techniques.

Fairing, blurring, blending. We outline categories of solid modeling techniques.

Tight hulls. We define the tight hull as a generalization of the convex hull that, like the relative convex hull, contains one set and excludes another. Unlike the relative convex hull, the tight hull minimizes total absolute Gaussian curvature, so it reflects properties of tight submanifolds. The tight hull also minimizes unsupported curvature, symmetrically behaving as a membrane in tension.

Tight blends. We define a tight hull separating the opening of an input shape from the opening of its complement as a tight blend. In two dimensions, tight blends resemble tightenings, and they offer the same curvature bound. We additionally explain the normal field simplification that proceeds as we vary the radius used for opening. In three dimensions, unsupported tight blends exhibit minimum Gaussian curvature rather than the saddles in unsupported tightening. Although the tight blend does not offer the same tightening curvature bound, it consistently simplifies normals.

Medial cover. We identify two-dimensional tight covers as sets that, like tight hulls, separate one set from another with rubber band-like boundary loops. After showing that an input shape composed of multiple components may admit an exponential number of tight covers, we define the medial cover by building sets from convex hulls of the contact points of maximal balls. We demonstrate that the medial cover is symmetric with respect to the sets the tight hull includes and excludes.

Implementation. Focusing on two-dimensional polygonal input, we implement tight covers by constructing the bounding loops of the tight hull by path planning in triangles lying in the gap between the triangles with vertices lying exclusively on sets inside the hull or lying exclusively on sets outside of the hull. Inserting contact points from maximal balls into the input, we use constrained Delaunay triangulation to obtain a medial over. Opening the input and its complement with a regular polygon, we approximate tight blends.

Applications. We explain the significance of tight hulls in convergent boundary estimation, which in particular enables accurate shading of images extracted from biomedical image data. We discuss the application of tight hulls to design with developable patches as well as the application of tight blends to normal field simplification. We conclude by indicating the tight hull's potential for polygonal repair.

11.3. Discussion

Our work involves blending and tightening boundaries of input shapes subject to the constraint that the output shape includes a subset of the input and excludes a subset of the input's complement. Blending replaces thin, sharp features with thick, rounded features, while tightening hugs the input's constraints by minimizing the output's boundary measure.

The techniques we use involve mathematical morphology, fairing, and convexity. In our work, morphology manipulates sets of balls, consequently capable of blending, defining constraints, enabling analysis of regularity, and transforming an input set to a nearly regular output. Fairing can optimize local curvatures and minimize boundary measure by deforming a surface with a flow defined by a set of scaled normal vectors. Given constraints established by morphology, fairing provides both blending and tightening. Convexity similarly hugs constraints while introducing simple patches of unsupported boundary.

Our work involving morphology establishes the significance of symmetry. Asymmetric operators tend to globally add or reduce material. As examples, growing and morphological closing add material, while shrinking and morphological opening remove it. Mason directly reduces the measure of changed material, while tightening, the tight hull, and tight blends avoid bias. We characterize avoiding bias as a form of fidelity to the input shape. The significance of fidelity is that symmetry treats features defined by positive space as though they have the same properties as features defined by negative space.

We argue that the tight hull, a symmetric operator, offers a particularly important technique. By generalizing the convex hull, the tight hull reflects set-theoretic aspects of the relative convex hull as well as morphological operators such as rounding and closing. By minimizing slack and unsupported slack, the tight hull also reflects the mean curvature flow used for tightening. Like mason and tightening, the tight hull is symmetric with respect to the input sets inside it and outside of it. In two dimensions, tight blends establish both a curvature bound and normal field simplification. Tight hulls also establish topological variety.

The tight hull's combination of properties offers flexibility and simplicity. Simplicity and flexibility are closely related, because simplicity is precisely what makes the tight hull generally applicable. Minimizing first Gaussian curvature and then unsupported curvature produces a boundary that is generally faithful to the input, including its geometry, its normal field, its interaction with light, and the balance between its interior and exterior. The tight hull also promises simplicity in storage, rendering, and manufacture.

Our central problem involves boundary simplification with limited tolerance. Using the mortar, we can construct nearly regular sets that are locally thick and smooth, with bounded error introduced by sampling [**2, 62**]. Without the mortar, we can remove errors from input models and reconstruct from samples [**24**]. Just as the convex hull is broadly applicable, our techniques address these basic aspects of solid modeling.

11.4. Future work

We anticipate four lines of future work. The first is to develop a body of theory concerning the tight hull. The second is to implement the tight hull for three-dimensional polyhedra, targeting an efficient, precise algorithm. We would also like to develop

tight blends, both to compare their properties to other blends and to consider blending with variable radii. Finally, we would like to define a regularizing operator, as we know of no regularizing technique with attractive properties.

In our work to develop the tight hull, we may benefit by employing existing tools [**9, 10, 13, 14, 27, 28, 37, 56**]. For instance, this may facilitate extension of our slack definition (Section A.3.) We also may develop an essentially equivalent definition of the tight hull, i.e. minimize unsupported slack before minimizing slack. Alternative definitions may exhibit a pattern of development comparable to those of various definitions of the medial axis. As definitions and theorems have grown out of the medial axis, we may likewise cultivate tools relevant to the tight hull. At the least, we would like to prove the tight hull's existence while extending its definition to a larger domain.

Development of the tight hull may involve incorporation of Gaussian flow techniques. One body of work has addressed the deformation of rolling stones. Early work established that rolling convex bodies shrink to spheres [**1, 29**], where points with higher Gaussian curvature have a greater rate of impact and higher volume loss than points with lower Gaussian curvature. Later work accomodated nonconvex bodies [**38**], restricting deformation to the convex subset of the boundary.

Deformation in $\mathbb{R}^3$ exhibits the problems generated by Gaussian flow that motivate restriction of the flow to convex points [**79**]. Displacing a point by an inward-pointing normal scaled by Gaussian curvature shifts both convex and concave points inward, but Gaussian curvature increases at concave points while decreasing at convex points. Increasing Gaussian curvature introduces instability, which leads Zhao and Xu [**79**] to displace concave points outward and appropriately handle hyperbolic points. We speculate that other forms of Gaussian flow are possible, such as decreasing the magnitude of the principal curvature of least magnitude.

We can also characterize Gaussian flow as a means of solving a Monge-Ampère equation [**71**], where with the tight hull we would seek zero Gaussian curvature. Caffarelli has obtained relevant regularity properties [**11, 12**]. We do not anticipate Gaussian flow yielding a significant change in the definition of the tight hull, but it may play a role in proofs and implementations.

As indicated by Chazelle's work [**17**], a tight hull with a polygonal three-dimensional boundary representation may have $O(n^2)$ introduced vertices given n input vertices. Computing the tight hull consequently requires at least quadratic worst-case time in that case. We tentatively suspect that the tight hull for triangle mesh input can be computed in low-order polygonal time, and we also consider pixel and voxel representations.

In our efforts to compare tight blending to other smoothing techniques, we anticipate emphasizing the difference between fairing and set theoretic approaches such as mathematical morphology and convexity. The tight blend is in some sense a bridge between the two, because the mortar is set-theoretic while the tight hull is defined by integral boundary properties.

We can compare tight blending to variable-radius blending [**73**] as two particular forms of smoothing. Varible-radius blending uses a control surface to locally determine the radii of balls that perform forms of opening and closing. Although variable-radius blending is set-theoretic, its output qualitatively resembles fairing because we can manipulate the control surface to obtain an aesthetically simple, appealing output boundary. We can define a tight hull as a set that contains the intersection of two sets - an input model and the interior of a control surface - while excluding the intersections of their complements. This defines a set of balls lying in the gap between the contained and excluded sets. The radii of these balls may approach zero, so the tight hull defined by the tight-radius blend does not guarantee the same kind of curvature bound exhibited by constant-radius tightenings and tight blends in two dimensions. Potentially extending the tight blend to the use of variable-radius balls

in mathematical morphology apparently emphasizes set theory, but it presumably involves a technique for modulating radii.

Our work on mason and tightening originally expressed efforts toward creating a regularizing operator, which would map an input shape to a thick, round output with bounded principal curvatures. We can demonstrate that it is impossible to regularize a shape with changes confined to its mortar, which raises the issue of how closely the output geometry can match the input. Our interest in regularization stems from the fact that volumetrically sampled regular shapes can be reconstructed with limited error and no topological change.

Broadly, we expect to develop combinations of morphology, fairing, and convexity.

APPENDIX A

Tight Hull Formalization

A.1. Polyhedra

A.1.1. Polyhedral solid. A set $S \subseteq \mathbb{R}^d$ is a *polyhedral solid* if and only if $S = S^{\circ-}$ and ∂S is a topological manifold contained in the union of a finite number of halfspace boundaries.

A.2. Tightening

A.2.1. Unique tightening. $\mathrm{T}_r(S)$ denotes a unique tightening of a polyhedral solid $S \subseteq \mathbb{R}^d$ with radius $r \in \mathbb{R}^>$.

A.3. Slack

A.3.1. Slack vectors. The set of *slack vectors* for polyhedral solid $S \subseteq \mathbb{R}^d$ at point $p \in \partial S$, denoted $\mathrm{SV}^S(p)$, is defined as

$$\begin{aligned}\mathrm{SV}^S(p) := \big\{ v \in \mathbb{S}^{d-1} \mid \exists r \in \mathbb{R}^> : \forall \rho \in (0,r) : \exists \sigma \in \mathbb{R}^> : \forall \tau \in (0,\sigma) : \exists \mathrm{T}_\tau(S) : \\ \exists \delta \in \mathbb{R}^> : \forall \epsilon \in (0,\delta) : \exists q \in (B_\rho(p) \cap \partial \mathrm{T}_\tau(S)) : q + \epsilon v \notin \mathrm{T}_\tau(S) \uparrow_\epsilon \big\}\end{aligned}$$

A.3.2. Slack count. The *slack count* for polyhedral solid $S \subseteq \mathbb{R}^d$ at direction $v \in \mathbb{S}^{d-1}$, denoted $\mathrm{SC}^S(v)$, is defined as

$$\mathrm{SC}^S(v) := \#\kappa\left(\left\{ p \in \partial \mathrm{T}_\sigma(S) \,\middle|\, v \in \mathrm{SV}^{\mathrm{T}_\sigma(S)}(p) \right\}\right)$$

Where

$$\exists \sigma \in \mathbb{R}^{>} : \exists \mathrm{T}_{\sigma}\left(S\right) : \forall \tau \in \left(0, \sigma\right) : \exists \mathrm{T}_{\tau}\left(S\right) :$$
$$\#\kappa\left(\left\{p \in \partial \mathrm{T}_{\tau}\left(S\right) \middle| v \in \mathrm{SV}^{\mathrm{T}_{\tau}(S)}\left(p\right)\right\}\right) = \#\kappa\left(\left\{p \in \partial \mathrm{T}_{\sigma}\left(S\right) \middle| v \in \mathrm{SV}^{\mathrm{T}_{\sigma}(S)}\left(p\right)\right\}\right)$$

A.3.3. Slack. The *slack* for polyhedral solid $S \subseteq \mathbb{R}^{d}$, denoted SL^{S}, is defined as

$$\mathrm{SL}^{S} := \sum_{i=1}^{\infty} \mathrm{H}^{d-1}\left(\left\{v \in \mathbb{S}^{d-1} \middle| \mathrm{SC}^{S}\left(v\right) \geq i\right\}\right)$$

A.4. Convexity

A.4.1. Two-dimensional section. The *two-dimensional section* of polyhedral solid $S \subseteq \mathbb{R}^{d}$ through $p \in \partial S$ with $v \in \mathrm{SV}^{S}\left(p\right)$ and $w \in \left\{x \in \mathbb{S}^{d-1} \mid v \cdot x = 0\right\}$ is denoted $\Omega^{S}\left(p, v, w\right)$ and defined as

$$\Omega^{S}\left(p, v, w\right) := \mathrm{AH}\left(\left\{p, p + v, p + w\right\}\right)$$

A.4.2. Sectional convexity. The *sectional convexity* of two-dimensional section $\Omega^{S}\left(p, v, w\right)$ is denoted $\sigma\left(\Omega^{S}\left(p, v, w\right)\right)$ and defined as

$$\sigma\left(\Omega^{S}\left(p, v, w\right)\right) := \begin{cases} 1 & \exists \delta \in \mathbb{R}^{>} : \forall \epsilon \in \left(0, \delta\right) : p + \epsilon v \notin \left(S \cap \Omega^{S}\left(p, v, w\right)\right) \uparrow_{\epsilon} \\ 0 & \nexists \delta \in \mathbb{R}^{>} : \forall \epsilon \in \left(0, \delta\right) : p + \epsilon v \notin \left(S \cap \Omega^{S}\left(p, v, w\right)\right) \uparrow_{\epsilon} \end{cases}$$

A.4.3. Convexity vector. The *convexity vector* of polyhedral solid $S \subseteq \mathbb{R}^{d}$ given $p \in \partial S$ and $v \in \mathrm{SV}^{S}\left(p\right)$, denoted $\mathrm{CV}^{S}\left(p, v\right)$, is defined as

$$\mathrm{CV}^{S}\left(p, v\right) := \left[\frac{\mathrm{H}^{d-2}\left(\left\{w \in \mathbb{S}^{d-1} \middle| \left(v \cdot w = 0\right) \wedge \left(\sigma\left(\Omega^{S}\left(p, v, w\right)\right) = 1\right)\right\}\right)}{\mathrm{H}^{d-2}\left(\mathbb{S}^{d-2}\right)}\right] v$$

A.5. Support

A.5.1. Support vector. Given polyhedral solids $S \subseteq \mathbb{R}^d$, $R \subseteq \mathbb{R}^d$, and $G \subseteq \mathbb{R}^d$ such that $R \subseteq S$ and $S \subseteq G^{c-}$, the *support vector* for S containing R relative to G given $p \in \partial S$ and $v \in \mathrm{SV}^S(p)$, denoted $\Psi^S_{R|G}(p, v)$, is defined as

$$\Psi^S_{R|G}(p, v) := \left[\frac{\mathrm{H}^{d-2}\left(\left\{ w \in \mathbb{S}^{d-1} \,|(v \cdot w = 0) \wedge (\alpha \vee \beta) \right\}\right)}{\mathrm{H}^{d-2}\left(\mathbb{S}^{d-2}\right)} \right] v$$

Where α denotes

$$\left(\sigma\left(\Omega^S(p, v, w)\right) = 1\right) \wedge \left(\sigma\left(\Omega^R(p, v, w)\right) = 1\right)$$

And β denotes

$$\left(\sigma\left(\Omega^{S^{c-}}(p, v, w)\right) = 1\right) \wedge \left(\sigma\left(\Omega^G(p, v, w)\right) = 1\right)$$

A.5.2. Support and patch support. Given polyhedral solids $S \subseteq \mathbb{R}^d$, $R \subseteq \mathbb{R}^d$, and $G \subseteq \mathbb{R}^d$ such that $R \subseteq S$ and $S \subseteq G^{c-}$, the *support* for S containing R relative to G at direction $v \in \mathbb{S}^{d-1}$, denoted $\Sigma^S_{R|G}(v)$, is defined as

$$\Sigma^S_{R|G}(v) := \sum_{\omega \in \kappa\left(\left\{ p \in \partial S \middle| v \in \mathrm{SV}^S(p) \right\}\right)} \sum_{\nu \in \left\{ \Psi^S_{R|G}(p,v) \,|\, p \in \omega \right\}} \frac{\|\nu\| \, \mathrm{H}^{\mathrm{D}(\omega)}\left(\Pi^S_{R|G}(\omega, \nu)\right)}{\mathrm{H}^{\mathrm{D}(\omega)}(\omega)}$$

Where *patch support* for S containing R relative to G on patch ω with support vector ν, denoted $\Pi^S_{R|G}(\omega, \nu)$, is defined as

$$\Pi^S_{R|G}(\omega, \nu) := \left\{ p \in \omega \mid \Psi^S_{R|G}(p, v) = \nu \right\}$$

While $\omega \in \kappa\left(\left\{ p \in \partial S \,\middle|\, v \in \mathrm{SV}^S(p) \right\}\right)$.

A.5.3. Unsupported slack. The *unsupported slack* for polyhedral solid $S \subseteq \mathbb{R}^d$, denoted $\Upsilon^S_{R|G}$, is defined as

$$\Upsilon^S_{R|G} := \mathrm{SL}^S - \int_{\mathbb{S}^{d-1}} \Sigma^S_{R|G}(v)\, dv$$

A.6. Tight hull

A.6.1. Inclusion/exclusion hulls. Given polyhedral solids $R \subseteq \mathbb{R}^d$ and $G \subseteq \mathbb{R}^d$ such that R and G^c are bounded, let N be the set of *inclusion/exclusion hulls* containing every polyhedral solid $X \subseteq \mathbb{R}^d$ such that $R \subseteq X$ and $X \subseteq G^{c-}$.

A.6.2. Candidate hulls. Given the set N of inclusion/exclusion hulls defined by polyhedral solids $R \subseteq \mathbb{R}^d$ and $G \subseteq \mathbb{R}^d$, let C be the set of *candidate hulls* containing every $X \in N$ such that there is no $Y \in N$ such that $\mathrm{SL}^Y < \mathrm{SL}^X$.

A.6.3. Tight hulls. Given the set C of candidate hulls defined by polyhedral solids $R \subseteq \mathbb{R}^d$ and $G \subseteq \mathbb{R}^d$, let T be the set of *tight hulls* containing every $X \in C$ such that there is no $Y \in C$ such that $\Upsilon^Y_{R|G} < \Upsilon^X_{R|G}$.

A.6.4. Tight hull. Given the set T of tight hulls defined by polyhedral solids $R \subseteq \mathbb{R}^d$ and $G \subseteq \mathbb{R}^d$, a *tight hull* of R relative to G, denoted $\mathrm{TH}(R|G)$, is an element of T.

A.7. Tight hull properties

A.7.1. Nonuniqueness. Given polyhedral solids $R \subseteq \mathbb{R}^d$ and $G \subseteq \mathbb{R}^d$, we can prove by example that $\mathrm{TH}(R|G)$ may be nonunique (Figure 6.21.)

A.7.2. Existence. Given that the set N of inclusion/exclusion hulls defined by polyhedral solids $R \subseteq \mathbb{R}^d$ and $G \subseteq \mathbb{R}^d$ is nonempty, we conjecture that a tight hull $\mathrm{TH}(R|G)$ exists.

A.7.3. Complementary. Given polyhedral solids $R \subseteq \mathbb{R}^d$ and $G \subseteq \mathbb{R}^d$, we conjecture that if G^c is bounded and $\mathrm{TH}(R|G)$ exists, then there exists a $\mathrm{TH}(G|R)$ such that $\mathrm{TH}(R|G) = \mathrm{TH}(G|R)^{c-}$ (Figure 6.29-6.32.)

A.7.4. Convex hull generalization.

Theorem A.1. *Given polyhedral solids $R \subseteq \mathbb{R}^d$ and $G \subseteq \mathbb{R}^d$ such that G^c is bounded and $\mathrm{CH}(R) \subseteq G^c$, there exists a unique $\mathrm{TH}(R|G)$ and $\mathrm{TH}(R|G) = \mathrm{CH}(R)$.*

Proof. The elements of inclusion/exclusion hulls N contain R, and the set C of candidate hulls defined by N consists of bounded convex sets because they minimize slack with the minimum possible slack value of $\mathrm{H}^{d-1}\left(\mathbb{S}^{d-1}\right)$ for polyhedral solids in G^c. Then $\mathrm{CH}(R) \in C$, because $\mathrm{CH}(R)$ is a bounded convex polyhedral solid containing R. Every vertex of $\mathrm{CH}(R)$ coincides with a convex vertex of R, so $\mathrm{CH}(R)$ has zero unsupported slack. If $x \in C$ and $\mathrm{CH}(R) \setminus x \neq \emptyset$ then x is nonconvex, which yields a contradiction. If $x \setminus \mathrm{CH}(R) \neq \emptyset$ then x has an unsupported vertex, so x has nonzero unsupported slack. Consequently, $\mathrm{TH}(R|G)$ is unique and $\mathrm{TH}(R|G) = \mathrm{CH}(R)$ (Figure 6.20.) □

REFERENCES

[1] ANDREWS, B. 1991. Gauss curvature flow: The fate of the rolling stones. *Inventiones Mathematicae 138*, 151-161.

[2] 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.

[3] 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.

[4] 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.

[5] BERTRAND, J., DIQUET, C.-F., AND PUISEUX, V. 1848. Démonstration d'un théorème de Gauss. *Journal de Mathématiques Pures et Appliquées 13*, 80-90.

[6] 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.

[7] BOBENKO, A., AND SCHRÖDER, P. 2005. Discrete Willmore flow. *Proceedings of the 3rd Eurographics Symposium on Geometry Processing*, Vienna, Austria, July 2005, M. DESBRUN AND H. POTTMANN, Eds. Eurographics Association, Aire-la-Ville, Switzerland, 101-110.

[8] BREHM, U. AND KÜHNEL, W. 1982. Smooth approximation of polyhedral surfaces regarding curvatures. *Geometriae Dedicata 12*, 61-85.

[9] BRÖCKER, L. AND KUPPE, M. 2000. Integral geometry of tame sets. *Geometriae Dedicata 82*, 285-323.

[10] CAFFARELLI, L. 1982. The two-obstacle problem for the biharmonic operator. *Pacific Journal of Mathematics 103*, 325–335.

[11] CAFFARELLI, L. 1990. Interior regularity of solutions to Monge-Ampère equations. *Contemporary Mathematics 107*, 13-17.

[12] CAFFARELLI, L. 1991. Some regularity properties of solutions of Monge Ampère equation. *Communications on Pure and Applied Mathematics 44*, 965-969.

[13] CAFFARELLI, L. 1998. *The Obstacle Problem*. Scuola Normale Superiore, Pisa, Italy.

[14] CAFFARELLI, L. AND MCCANN, R. 2010. Free boundaries in optimal transport and Monge-Ampère obstacle problems. *Annals of Mathematics 171*, 673-730.

[15] 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.

[16] 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.

[17] CHAZELLE, B. 1984. Convex partitions of polyhedra: A lower bound and worst-case optimal algorithm. *SIAM Journal of Computing 13*, 488-507.

[18] CHERN, S.-S. AND LASHOF, R. 1957. On the total curvature of immersed manifolds. *American Journal of Mathematics 79*, 306-318.

[19] CHEW, L. 1987. Constrained Delaunay triangulations. *Proceedings of the 3rd Annual Symposium on Computational Geometry*, Waterloo, Canada, June 1987, ACM Press, 215-222.

[20] CHOI, H., CHOI, S., AND MOON, H. 1997. Mathematical theory of medial axis transform. *Pacific Journal of Mathematics 181*, 57-88.

[21] CHOPP, D. 1993. Computing minimal surfaces via level set curvature flow. *Journal of Computational Physics 106*, 77-91.

[22] DANIELSSON, P. 1980. Euclidean distance mapping. *Computer Graphics and Image Processing 14*, 227-248.

[23] 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.

[24] DEY, T. 2007. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge University Press, New York, NY.

[25] EVANS, L. AND SPRUCK, J. 1991. Motion of level sets by mean curvature I. *Journal of Differential Geometry 33*, 635-681.

[26] FAIRY, S. 2003. http://www.cytology-asc.com/branch/nsw/jun03/cases.html.

[27] FEDERER, H. 1959. Curvature measures. *Transactions of the American Mathematical Society 93*, 418-491.

[28] FEDERER, H. 1969. *Geometric Measure Theory*. Springer-Verlag, New York, NY.

[29] FIREY, W. 1974. Shapes of worn stones. *Mathematika 21*, 1-11.

[30] FREY, W. 2002. Boundary triangulations approximating developable surfaces that interpolate a space curve. *International Journal of Foundations of Computer Science 13*, 285-302.

[31] GORECKA-EGAN, E. Helen of Troy. In *Paper Sculpture*, by MCPHARLIN, P., 1944, Marquardt & Company, New York, NY, pp. 46.

[32] GRAYSON, M. 1987. The heat equation shrinks embedded plane curves to round points. *Journal of Differential Geometry 26*, 285-314.

[33] HAUSDORFF, F. 1918. Dimension und äusseres Mass. *Mathematische Annalen 79*, 157-179.

[34] HE,T., HONG, L., VARSHNEY, A., AND WANG, S. 1996. Controlled Topology Simplification. *IEEE Transactions on Visualization and Computer Graphics 2*, 171-184.

[35] HEIJMANS, H. 1994. *Morphological Image Operators.* Academic Press, Boston, MA.

[36] HERSHBERGER, J. AND SNOEYINK, J. 1994. Computing minimum length paths of a given homotopy class. *Computational Geometry: Theory and Applications 4*, 63-97.

[37] HINTERMÜLLER, M., KOVTUNENKO, V.A., AND KUNISCH, K. Obstacle problems with cohesion: A hemivariational inequality approach and its efficient numerical solution. *SIAM Journal on Optimization 21*, 491-516.

[38] ISHII, H. AND MIKAMI, T. 2001. A mathematical model of the wearing process of a nonconvex stone. *SIAM Journal on Mathematical Analysis 33*, 860-876.

[39] JORDAN, C. 1892. Remarques sur les intégrales définies. *Journal de Mathématiques Pures et Appliquées 8,* 69-99.

[40] KADUSHIN, R. 2006. Square dance coffee table. `http://ronen-kadushin.com/Open_Design.asp`.

[41] KILIAN, M., FÖRY, S., CHEN, Z., MITRA, N., SHEFFER, A., AND POTTMANN, H. 2008. Curved folding. *ACM Transactions on Graphics 27: SIGGRAPH 2008*, 75.

[42] KIM, B. AND ROSSIGNAC, J. 2005. GeoFilter: Geometric selection of mesh filter parameters. *Computer Graphics Forum 24: Eurographics 2005*, 295-302.

[43] KLETTE, R. AND ROSENFELD, A. 2004. *Digital Geometry: Geometric Methods for Digital Picture Analysis*. Morgan Kaufmann, San Francisco, CA.

[44] KÜHNEL, W. 2006. *Student Mathematical Library Volume 16: Differential Geometry: Curves — Surfaces — Manifolds*. American Mathematical Society, Providence, RI.

[45] KUIPER, N. 1985. Minimal total curvature, a look at old and new results. *Russian Mathematical Surveys 40*, 49-55.

[46] 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.

[47] LEE, D. AND PREPARATA, F. 1984. Euclidean shortest paths in the presence of rectilinear barriers. *Networks 14*, 393-410.

[48] 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.

[49] MORETON, H. 1992. *Minimum Curvature Variation Curves, Networks, and Surfaces for Free-Form Shape Design*. Doctoral dissertation, University of California at Berkeley, Berkeley, CA.

[50] 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.

[51] 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.

[52] PEANO, G. 1887. *Applicazioni Geometriche del Calcolo Infinitesimale*. Fratelli Bocca Editori, Turin, Italy.

[53] 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.

[54] POTTMANN, H., ASPERL, A., HOFER, M., AND KILIAN, A. 2007. *Architectural Geometry*. Bentley Institute Press, Exston, PA.

[55] 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.

[56] RATAJ, J. AND ZÄHLE, M. 2003. Normal cycles of Lipschitz manifolds by approximation with parallel sets. *Differential Geometry and its Applications 19*, 113-126.

[57] 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.

[58] ROSE, N. 1985. *Mathematical Maxims and Minims*. Rome Press, Raleigh, NC.

[59] ROSSIGNAC, J. AND REQUICHA, A. 1984. Constant-radius blending in solid modeling. *ASME Computers in Mechanical Engineering 3*, 65-73.

[60] ROSSIGNAC, J. 1985. *Blending and Offsetting Solid Models*. Doctoral dissertation, University of Rochester, Rochester, NY.

[61] SCARINCI, I. 2008. Financial scandal at the Guggenheim Bilbao. Translated by MAHABIR, A. *ARCADJA artMagazine*, `http://www.arcadja.com`.

[62] SERRA, J. 1982. *Image Analysis and Mathematical Morphology, Volume 1*. Academic Press, New York, NY.

[63] 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.

[64] SKLANSKY, J. AND KIBLER, D. 1976. A theory of nonuniformly digitized pictures. *IEEE Transactions on Systems, Man, and Cybernetics 6*, 637-647.

[65] 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.

[66] STENGER, R. 2002. http://www.cnn.com/2002/TECH/space/10/25/galactic.crash.

[67] STRUIK, D. 1961. *Lectures on Classical Differential Geometry*. Dover Publications, New York, NY.

[68] TAUBIN, G. 1995. Curve and surface smoothing without shrinkage. *Proceedings of the 5th Annual Conference on Compuer Vision: ICCV 95*, 852.

[69] 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.

[70] TOUSSAINT, G. 1989. Computing geodesic properties inside a simple polygon. *Revue D'Intelligence Artificielle 3*, 9-42.

[71] URBAS, J. 1997. Self-similar solutions of Gauss curvature flows. *Proceedings of the CBMS Conference on the Monge-Ampère Equation: Applications to Geometry and Optimization*, Deerfield Beach, FL, July 1997, Florida Atlantic University, Broward Campus.

[72] VAN DE VEL, M. 1993. *North-Holland Mathematical Library Volume 50: Theory of Convex Structures*. North-Holland, New York, NY

[73] WHITED, B., AND ROSSIGNAC, J. 2009. Relative blending. *Computer-Aided Design 41*, 456-462.

[74] WILLIAMS, J. AND ROSSIGNAC, J. 2005. Mason: Morphological simplification. *Graphical Models 67*, 285-303.

[75] WILLIAMS, J. AND ROSSIGNAC, J. 2007. Tightening: Morphological simplification. *International Journal of Computational Geometry and Applications 17*, 487-503.

[76] WILLMORE, T. 1971. Tight immersions and total absolute curvature. *Bulletin of the London Mathematical Society 3*, 129-151.

[77] 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.

[78] 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.

[79] ZHAO, H. AND XU, G. 2006. Triangular surface mesh fairing via Gaussian curvature flow. *Journal of Compuation and Applied Mathematics 195*, 1-12.