# SELF-DELAUNAY MESHES FOR SURFACES 

by<br>Ramsay Dyer<br>B.Sc., University of British Columbia, 1990<br>M.Sc., Université de Montréal, 1996<br>a Thesis submitted in Partial fulfillment OF THE REQUIREMENTS FOR THE DEGREE OF Doctor of Philosophy<br>in the School<br>of<br>Computing Science<br>(C) Ramsay Dyer 2010<br>SIMON FRASER UNIVERSITY<br>Spring 2010

All rights reserved. This work may not be reproduced in whole or in part, by photocopy or other means, without the permission of the author.

## APPROVAL

| Name: | Ramsay Dyer |
| :--- | :--- |
| Degree: | Doctor of Philosophy |
| Title of Thesis: | Self-Delaunay meshes for surfaces |

## Examining Committee: Dr. Mirza Faisal Beg

Chair

Dr. Torsten Möller, Senior Supervisor

Dr. Hao Zhang, Senior Supervisor

Dr. Binay Bhattacharya, SFU Examiner

Dr. Leonidas J. Guibas, External Examiner, Professor of Computer Science, Stanford University

## Abstract

In the Euclidean plane, a Delaunay triangulation can be characterized by the requirement that the circumcircle of each triangle be empty of vertices of all other triangles. For triangulating a surface $S$ in $\mathbb{R}^{3}$, the Delaunay paradigm has typically been employed in the form of the restricted Delaunay triangulation, where the empty circumcircle property is defined by using the Euclidean metric in $\mathbb{R}^{3}$ to measure distances on the surface. More recently, the intrinsic (geodesic) metric of $S$ has also been employed to define the Delaunay condition. In either case the resulting mesh $M$ is known to approximate $S$ with increasing accuracy as the density of the sample points increases. However, the use of the reference surface $S$ to define the Delaunay criterion is a serious limitation. In particular, in the absence of the original reference surface, there is no way of verifying if a given mesh meets the criterion.

We define a self-Delaunay mesh as a triangle mesh that is a Delaunay triangulation of its vertex set with respect to the intrinsic metric of the mesh itself. This yields a discrete surface representation criterion that can be validated by the properties of the mesh alone, independent of any reference surface the mesh is supposed to represent. The intrinsic Delaunay triangulation that characterizes self-Delaunay meshes makes them a natural domain for discrete differential geometry, and the discrete exterior calculus in particular.

We examine self-Delaunay meshes and their relationship with other Delaunay structures for surface representation. We study sampling conditions relevant to the intrinsic approach, and compare these with traditional sampling conditions which are based on extrinsic quantities and distances in the ambient Euclidean space. We also provide practical and provably correct algorithms for constructing self-Delaunay meshes. Of particular interest in this context is the extrinsic edge flipping algorithm which extends the familiar algorithm for producing planar Delaunay triangulations.

Keywords: triangle mesh; surface meshing; Delaunay triangulation; Delaunay edge flip

## Acknowledgements

I would like to thank Brian Funt for encouraging me to pursue graduate studies in computing science.

My supervisors, Torsten Möller and Richard Zhang, have been a reliable source of encouragement, guidance, and insight. I am very grateful to them for their confidence in me and for allowing me the freedom I needed to establish my footing.

Binay Bhattacharya and Leonidas Guibas provided insightful suggestions which have significantly improved this thesis. I really appreciate this input. In particular, I want to thank Professor Guibas for encouraging me to find a more distinctive name for the principle object of study.

My experience at the GrUVi Lab at SFU has been enjoyable and rewarding. I have received a tremendous benefit from interactions with the other students in the lab; I would like to thank them all. In particular, I owe a debt to John Li who provided the implementation of his non-obtuse decimation algorithm, which I modified to produce self-Delaunay meshes.

Finally, I am grateful for the support from my family and friends. I cannot thank Cheryl enough for the unyielding love and support she has bestowed upon me.

## Contents

Approval ..... ii
Abstract ..... iii
Acknowledgements ..... iv
Contents ..... V
List of Figures ..... ix
List of Symbols ..... xii
Preface ..... xv
1 Introduction ..... 1
2 Background ..... 7
2.1 The Euclidean setting ..... 8
2.1.1 Voronoi diagrams ..... 8
2.1.2 Delaunay triangulations ..... 10
2.1.3 Generalizations and related structures ..... 18
2.2 Triangulations and triangle meshes for surfaces ..... 21
2.2.1 Embeddings and Immersions ..... 21
2.2.2 Surfaces ..... 23
2.2.3 Surface triangulations ..... 26
2.2.4 Triangle meshes ..... 26
2.3 Delaunay structures for surfaces ..... 29
2.3.1 Voronoi diagrams on surfaces ..... 30
2.3.2 Restricted Delaunay triangulation (rDt) ..... 33
2.3.3 Intrinsic Delaunay triangulation mesh (iDt-mesh) ..... 35
2.3.4 Related structures ..... 37
2.4 Surface approximation theory ..... 44
2.4.1 Sampling and topological consistency ..... 45
2.4.2 Convergence and geometric accuracy ..... 49
2.4.3 Element quality ..... 52
3 Self-Delaunay meshes ..... 55
3.1 Delaunay and Voronoi structures on pwf Surfaces ..... 56
3.1.1 Delaunay Triangulations on pwf Surfaces ..... 58
3.1.2 Voronoi Diagrams on pwf Surfaces ..... 60
3.1.3 Proper Triangulations ..... 64
3.2 Self-Delaunay meshes ..... 66
3.2.1 Delaunay extrinsic edge flips ..... 68
3.2.2 Smooth self-Delaunay meshes ..... 69
3.3 Distinctions between the Delaunay structures ..... 71
3.3.1 The iDt-mesh is not a self-Delaunay mesh ..... 71
3.3.2 The rDt is not a self-Delaunay mesh ..... 72
3.4 Discussion ..... 73
3.4.1 A fourth Delaunay surface structure ..... 75
4 Gabriel meshes are self-Delaunay meshes ..... 76
4.1 Background ..... 77
4.2 Gabriel complexes ..... 78
4.2.1 Gabriel faces in a tetrahedron ..... 80
4.3 The anatomy of a hinge ..... 81
4.3.1 Equivalent hinges ..... 81
4.3.2 Cross-sectional diagrams ..... 82
4.4 Relating Gabriel and Delaunay properties ..... 84
4.4.1 Symmetric cases ..... 84
4.4.2 Asymmetric cases ..... 85
4.5 Obstructions to Gabriel meshes ..... 87
4.6 Proof of Theorem 4.1 ..... 88
4.7 Discussion ..... 91
5 Intrinsic surface sampling criteria ..... 94
5.1 Sizing functions for surface sampling ..... 94
5.1.1 A natural hierarchy of extrinsic sizing functions ..... 95
5.1.2 Strong convexity ..... 96
5.1.3 Intrinsic sizing functions ..... 97
5.2 Intrinsic sampling criteria ..... 100
5.2.1 The closed ball property via strong convexity ..... 100
5.2.2 Sampling via the strong convexity radius ..... 103
5.2.3 A weaker criterion: the intrinsic sampling radius ..... 103
5.3 Relating extrinsic and intrinsic sampling criteria ..... 106
5.3.1 A lfs estimate on the intrinsic sampling radius ..... 107
5.3.2 Bounding geodesic lengths ..... 111
5.3.3 Extrinsic criteria meeting intrinsic demands ..... 111
5.4 Discussion ..... 113
6 Constructing self-Delaunay meshes ..... 114
6.1 Geometry-preserving Delaunay remeshing ..... 115
6.1.1 Remeshing algorithm ..... 115
6.1.2 Termination proof ..... 117
6.1.3 Meshes with boundaries ..... 120
6.2 Delaunay remeshing via mesh edge flipping ..... 120
6.2.1 Edge flipping and refinement algorithm ..... 120
6.2.2 Delaunay extrinsic edge flipping and area minimization ..... 121
6.2.3 Unflippable edges ..... 123
6.3 Delaunay mesh decimation ..... 124
6.4 Experimental results ..... 130
6.5 Discussion ..... 134
7 Analysis of Delaunay extrinsic edge flips ..... 136
7.1 Non-sharp and smooth meshes ..... 138
7.2 Delaunay extrinsic edge flips and triangle circumradius ..... 141
7.3 Relation to the ambient Delaunay tetrahedralization ..... 145
7.4 Discussion ..... 148
8 Conclusions ..... 149
A Circumcentric dual cells with negative area ..... 153
A. 1 Circumcentric dual cells and their area ..... 153
A. 2 Construction of a dual cell with negative area ..... 156
A. 3 Implications for the DEC on arbitrary meshes ..... 158
B Sphere lemmas ..... 161
B. 1 Sphere intersections ..... 161
B. 2 Spherical geometry ..... 163
C More on the intrinsic Voronoi diagram ..... 167
C. 1 On the assumption of general position ..... 167
C. 2 Chavel's convexity lemma ..... 170
C.2.1 Overview ..... 170
C.2.2 An outline of the proof ..... 171
C. 3 Improving the bound in Theorem 5.19 ..... 174
D Allowable region calculations ..... 179
D. 1 The feasible region ..... 179
D.1.1 Feasible points on the horizontal plane ..... 180
D.1.2 The Lifting Lemma ..... 181
D.1.3 Feasible points on the circumcylinder ..... 183
D. 2 Linear allowable region boundaries ..... 185
D.2.1 Opposite-side case ..... 186
D.2.2 Same-side case ..... 189
Bibliography ..... 192
Index ..... 202

## List of Figures

1.1 The rDt and the iDt-mesh ..... 2
2.1 Voronoi diagram ..... 8
2.2 Simplicial complex ..... 11
2.3 Euclidean domain ..... 17
2.4 Gabriel graph ..... 21
2.5 Embedding and immersion ..... 23
2.6 Cone geometry ..... 25
2.7 Flat triangles ..... 26
2.8 Not immersed ..... 27
2.9 A covering and its nerve ..... 31
2.10 A closed ball violation ..... 32
2.11 The rDt and the rVd ..... 34
2.12 Restricted Voronoi diagram ..... 36
2.13 Sliver tetrahedron ..... 40
2.14 Schwarz torus ..... 46
2.15 Medial axis ..... 47
2.16 Trefoil knot ..... 48
3.1 Close empty disks ..... 61
3.2 Not well formed - Non-proper ..... 65
3.3 Flip-tet ..... 69
3.4 Bad self-Delaunay mesh ..... 70
3.5 iDt is not a self-Delaunay mesh ..... 71
3.6 rDt is not a self-Delaunay mesh ..... 73
3.7 Voronoi distortion ..... 74
4.1 Equivalent hinges ..... 82
4.2 Locally Gabriel implies locally Delaunay ..... 84
4.3 Single certificate case ..... 86
4.4 Obstruction to a Gabriel mesh ..... 88
4.5 Gabriel face of a simplex ..... 89
4.6 Circumradius comparison ..... 89
4.7 Asymmetric Delaunay condition ..... 91
4.8 No flip algorithm for self-rDt meshes ..... 92
5.1 lfs vs. local reach. ..... 95
5.2 Exponential map and injectivity ..... 99
5.3 Strong convexity and Voronoi cells ..... 101
5.4 Comparing sampling criteria ..... 112
6.1 Edge split ..... 115
6.2 Naive edge splitting ..... 116
6.3 Geometry-preserving Delaunay remeshing ..... 117
6.4 Edge splitting termination proof ..... 119
6.5 Measures not optimized ..... 121
6.6 Unflippable edges ..... 123
6.7 Edge collapse ..... 125
6.8 Subtending edge constraints ..... 126
6.9 Incident edge constraints ..... 127
6.10 Allowable regions for incident edges ..... 128
6.11 Comparison of Delaunay remeshing algorithms ..... 132
6.12 Vertex counts ..... 133
6.13 Decimation example ..... 134
6.14 Flip distortion ..... 135
7.1 Non-sharp, but not smooth. ..... 138
7.2 Smooth to unflippable ..... 142
7.3 A flip-tet not in $\mathcal{D}^{3}(P)$ ..... 146
A. 1 A negative dual edge ..... 155
A. 2 Dual edge area contribution ..... 156
A. 3 Dual cell with negative area ..... 157
B. 1 Intersecting spheres ..... 162
B. 2 Spherical triangle ..... 164
C. 1 Accumulating Voronoi vertices ..... 169
C. 2 Chavel's lemma ..... 172
D. 1 The feasible region in the plane ..... 180
D. 2 Lifting Lemma ..... 182
D. 3 Feasible regions on circumcylinder ..... 184
D. 4 Opposite-side case ..... 186
D. 5 Comparing tangents ..... 188
D. 6 Same-side case ..... 189
D. 7 Exit tangent ..... 190

## List of Symbols

$P$ sample point set ..... 8
$n \quad$ number of samples (or vertices): $n=|P|$ ..... 8
$\mathbb{R}^{d} \quad d$-dimensional Euclidean space ..... 8
$V(p)$ Euclidean Voronoi cell of $p$ in $\mathbb{R}^{d}$ ..... 8
$d_{\mathbb{R}^{d}}(p, q) \quad$ Euclidean distance between $p$ and $q$ in $\mathbb{R}^{d}$ ..... 9
$V(q ; Q) \quad$ Voronoi cell of q with respect to point set $Q$ ..... 10
$|V| \quad$ size of set $V$ (cardinality or volume, as appropriate) ..... 10
$\mathcal{K}$ simplicial complex ..... 12
$K \quad$ carrier of simplicial complex $\mathcal{K}$ ..... 12
$\partial D \quad$ boundary of set $D$ ..... 12
$[p, q] \quad$ edge between $p$ and $q$ (a 1-simplex) ..... 12
$\mathcal{D}^{d}(P) \quad$ Euclidean Delaunay triangulation of $P \subset \mathbb{R}^{d}$ ..... 12
$[a, b, c] \quad$ triangle (a 2-simplex) ..... 13
$P_{w} \quad$ weighted point set ..... 19
$\pi_{p}(x) \quad$ power distance of $x$ from weighted point $p$ ..... 19
$V_{w}(p) \quad$ power cell of p ..... 19
$\left.\phi\right|_{V} \quad$ restriction of $\phi$ to $V$ ..... 22
$S \quad$ smooth surface ..... 23
$\mathcal{S} \quad$ abstract Riemannian 2-manifold ..... 23
$\|\boldsymbol{v}\| \quad$ norm of vector $\boldsymbol{v}$ ..... 24
$\hat{\boldsymbol{n}}_{S}(p) \quad$ surface normal at $p$ ..... 24
$\kappa_{1}(p), \kappa_{2}(p)$ principle curvatures at $p$ ..... 24
$\kappa(p) \quad$ maximal curvature at p ..... 24
$G(p) \quad$ Gaussian curvature at $p$ ..... 24

| $d_{S}(p, q)$ | geodesic distance between $p, q \in S$. . . . . . . . . . . . . . . . . . . . 24 |
| :---: | :---: |
| $\mathcal{M}$ | pwf surface . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 24 |
| $\mathcal{K}_{M}$ | manifold simplicial complex associated with M . . . . . . . . . . 26 |
| $U(p)$ | umbrella at p . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 27 |
| $\mathcal{A}$ | abstract simplicial complex . . . . . . . . . . . . . . . . . . . . . . . 28 |
| $\mathcal{A}_{M}$ | abstract simplicial complex associated with M . . . . . . . . . . . . . 28 |
| $\{i, j, k\}$ | abstract simplex (or non-Euclidean triangle, etc.) . . . . . . . . . . . 28 |
| $\mathcal{V}(p)$ | intrinsic Voronoi cell of $p$ on S . . . . . . . . . . . . . . . . . . . . . . 30 |
| $\left.V\right\|_{S}(p)$ | restricted Voronoi cell of $p$ on $S$. . . . . . . . . . . . . . . . . . . . . 30 |
| $V(u p q)$ | Voronoi edge dual to Delaunay simplex $[u, p, q]$. . . . . . . . . . . . 34 |
| $B_{S}(c ; r)$ | open geodesic disk centred at $c \in S$, with radius $r \ldots \ldots$ |
| $\bar{A}$ | closure of set A . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 35 |
| $\hat{\boldsymbol{v}}_{p}$ | pole vector at $p$. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 38 |
| $\overrightarrow{p x}$ | vector from $p$ to $x$. . . . . . . . . . . . . . . . . . . . . . . . . . . . 38 |
| $\angle_{a}(\boldsymbol{a}, \boldsymbol{b})$ | $\min \{\angle(\boldsymbol{a}, \boldsymbol{b}), \angle(-\boldsymbol{a}, \boldsymbol{b})\} \ldots \ldots$ |
| $\mathbb{R}^{+}$ | positive reals . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 48 |
| $B_{\mathbb{R}^{3}}(c ; r)$ | open Euclidean ball of centred at $c \in \mathbb{R}^{3}$, with radius $r$. . . . . . . 48 |
| $\xi_{S}$ | orthogonal projection mapping onto S . . . . . . . . . . . . . . . . . 50 |
| $U_{\rho_{R}}$ | tubular neighbourhood (adaptive) . . . . . . . . . . . . . . . . . . . . 50 |
| $\hat{\boldsymbol{n}}_{t}$ | normal vector to triangle t . . . . . . . . . . . . . . . . . . . . . . . . 50 |
| $r_{t}$ | circumradius of triangle $t$. . . . . . . . . . . . . . . . . . . . . . . . . 54 |
| $(u, e, v)$ |  |
| $\left(t_{1}, t_{2}\right)$ | hinge described by adjacent triangles . . . . . . . . . . . . . . . . . . 68 |
| $\sigma_{k}$ | $k$-simplex . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 78 |
| $c_{\sigma}$ | circumcentre of simplex $\sigma$. . . . . . . . . . . . . . . . . . . . . . . 78 |
| $r_{\sigma}$ | circumradius of simplex $\sigma$. . . . . . . . . . . . . . . . . . . . . . 78 |
| $B_{\sigma}$ | diametric $d$-ball of simplex $\sigma \subset \mathbb{R}^{d}$. . . . . . . . . . . . . . . . . . . 78 |
| $\mathcal{G}^{k}(P)$ | $k$-Gabriel complex of P . . . . . . . . . . . . . . . . . . . . . . . . . 78 |
| $S_{\sigma}$ |  |
| aff $(\sigma)$ | affine hull of $\sigma$. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 88 |
| $\rho_{f}(x)$ | local feature size at $x \in S$. . . . . . . . . . . . . . . . . . . . . . . . 95 |
| $\rho_{R}(x)$ | local reach at $x \in S$. . . . . . . . . . . . . . . . . . . . . . . . . . . . 95 |
| $\rho_{\kappa}(x)$ | maximal curvature radius at $x \in S$. . . . . . . . . . . . . . . . . . . 96 |

$\rho_{G}(x) \quad$ Gaussian curvature radius at $x \in S$ ..... 96
$\rho_{\mathrm{sc}}(x) \quad$ strong convexity radius at $x \in S$ ..... 97
$T_{x} S \quad$ tangent plane of $S$ at $x$ ..... 98
$\rho_{i}(x) \quad$ injectivity radius at $x \in S$ ..... 98
$\rho_{m}(x) \quad$ intrinsic sampling radius at $x \in S$ ..... 100
$S V(p, e) \quad$ split vertex on edge $e$ relative to end $p$ ..... 116
$a b \quad$ the short circular arc between $a$ and $b$ ..... 128
$\operatorname{sh}(t)$ orthogonal shadow of $t$ ..... 138
$F \quad$ boundary of the feasible region ..... 180
$C \quad$ circumcircle of the fixed vertices (Appendix D) ..... 180

## Preface

The topic of self-Delaunay meshes evolved out of research in spectral mesh processing. I was studying the robustness of the spectrum of mesh Laplacian operators with respect to changes in the mesh connectivity. The cotan operator was the most promising first order operator in this respect, and I then became aware of the work of Bobenko and Springborn [BS05], which provided a strong argument to suggest that the cotan operator should be defined with respect to the intrinsic Delaunay triangulation of the mesh, rather than via the triangulation which describes the mesh faces.

I then became interested in meshes for which the triangulation describing the mesh faces was itself the intrinsic Delaunay triangulation of the mesh. Naturally, I called these objects Delaunay meshes. This term was used in published papers whose content forms the majority of this thesis. However, as is made apparent in Chapter 2, there are many distinct Delaunay structures used for surface representation. It became evident that the generic term "Delaunay mesh" was an unfortunate choice for attempting to distinguish a specific Delaunay structure. It is hoped that the new name which is adopted here, self-Delaunay mesh, will fit more comfortably in the still evolving nomenclature for Delaunay structures for surfaces.

The first algorithm devised for constructing self-Delaunay meshes was the obvious one: take an arbitrary triangle mesh and flip all edges that aren't locally Delaunay. The algorithm seemed to converge, but in my first description of the results I had a footnote mentioning that on coarse meshes we could occasionally encounter an edge which was not locally Delaunay, but which could not be flipped because the edge it would flip to already exists in the mesh. The majority of this thesis has been motivated by an attempt to come to terms with that footnote.

The body of the thesis includes material which was previously published in four separate
papers. The primary contributions of these papers are contained in chapters according to the following correspondence:

- Chapter 3: Introduces self-Delaunay meshes, and demonstrates that they are distinct from the restricted Delaunay triangulation, and the mesh obtained through an intrinsic Delaunay triangulation of a smooth surface. [DZM07b]
(http://doi.acm.org/10.1145/1236246.1236306 © 2007 Association for Computing Machinery, Inc. Reprinted by permission.)
- Chapter 4: Explores the relationship between Gabriel meshes and self-Delaunay meshes. [DZM09]
(http://doi.acm.org/10.1145/1629255.1629293 © 2009 Association for Computing Machinery, Inc. Reprinted by permission.)
- Chapter 5: Develops intrinsic Voronoi-based sampling criteria for surfaces, and relates these to traditional extrinsic criteria. [DZM08]
(Journal compilation © 2008 The Eurographics Association and Blackwell Publishing Ltd. Reproduced by kind permission of the Eurographics Association.)
- Chapter 6: Describes algorithms with correctness guarantees for constructing selfDelaunay meshes. [DZM07a]
(© Eurographics Association 2007, Reproduced by kind permission of the Eurographics Association.)


## Chapter 1

## Introduction

The Delaunay triangulation of a point set in a planar or higher dimensional Euclidean domain is a fundamental data structure in computational geometry [dBvKOS98], and it has been used extensively in scientific computing for meshing Euclidean domains [She97]. It defines a natural neighbour relation amongst the points, and its popularity can be attributed to the many attractive properties it possesses. The Delaunay triangulation of a point set $P \subset \mathbb{R}^{2}$ can be characterized by the property that the circumcircle of each triangle is empty of points from $P$.

This thesis develops the theory of self-Delaunay meshes as representatives of smooth surfaces. We reserve the term intrinsic to apply uniquely to properties and objects of a surface that can be defined in terms of the surface itself, without any reference to its embedding in the ambient space. Self-Delaunay meshes are manifold triangle meshes with a connectivity conforming to an intrinsic Delaunay triangulation of the vertex set. The intrinsic Delaunay triangulation of the vertices of a mesh is always well defined, and it optimizes the same energy functionals as are optimized by the planar Delaunay triangulations [BS07]. In a self-Delaunay mesh, the intrinsic Delaunay triangulation of the vertices is given by the mesh connectivity itself.

There are other ways in which the Delaunay paradigm has been used to define a triangle mesh structure for representing smooth surfaces. The most common of these is the restricted Delaunay triangulation (rDt), which was formally introduced by Edelsbrunner and Shah [ES94]. The restricted Delaunay triangulation of a set of sample points $P$ on a surface $S$ can be defined as the substructure of the three dimensional Delaunay triangulation of $P$ which consists of those Delaunay triangles whose vertices lie on a sphere centred on $S$ and


Figure 1.1: The Delaunay paradigm with different metrics. (a) In the restricted Delaunay triangulation of $P$ on $S$, every Delaunay triangle has a Euclidean circumsphere centred on $S$ which is empty of points from $P$. (b) In the iDt-mesh, each triangle is such that its vertices lie on the boundary of a geodesic disk on $S$ that contains no points from $P$.
containing no elements of $P$, as shown in Figure 1.1(a).
Recently the power of an intrinsic approach to surface approximation has begun to be recognized [DLYG06]. In this approach we consider the geodesic distance between points on $S$, defined by the length of the shortest curve on $S$ that connects the points in question. A natural triangle mesh structure which arises in this context is defined by means of the intrinsic Delaunay triangulation (iDt) of $P$ on $S$. The connectivity of this mesh is given by the requirement that the vertices of every triangle face define a geodesic circumdisk on $S$ that is empty of elements of $P$. In other words, the three vertices lie on the boundary of a region $D \subset S$ and every point on the boundary of $D$ has the same geodesic distance to some $c \in S$ that (the centre of $D$ ), and there are no vertices in the interior of $D$. See Figure 1.1(b). We refer to such a mesh as an iDt-mesh.

A self-Delaunay mesh is one in which each triangle face has an empty geodesic circumdisk on the mesh itself. This is the important distinction between a self-Delaunay mesh and an rDt or an iDt-mesh. A self-Delaunay mesh does not depend on a reference surface for its
definition. In the absence of the original surface $S$, it is not possible in general to establish whether or not a given mesh is an rDt or an iDt-mesh. By contrast, it is easy to check whether or not a mesh is a self-Delaunay mesh. After the traditional Delaunay structures are reviewed in Chapter 2, we formally introduce self-Delaunay meshes in Chapter 3 and demonstrate that they are distinct from the rDt and the iDt-mesh, regardless of how densely the vertex set $P$ is sampled on $S$.

Although the self-Delaunay mesh is defined in terms of its own intrinsic metric and is independent of a reference surface, it is not intrinsically defined. Rather, the definition demands a marriage of intrinsic and extrinsic triangulations. The edges of a mesh can be interpreted as an artifact of the embedding in $\mathbb{R}^{3}$ : there is no way to discern mesh edges from a study of the metric properties of the mesh. Thus no mesh structure can be intrinsically defined. A self-Delaunay mesh demands that the intrinsic Delaunay triangulation coincides with this extrinsic triangulation that defines the mesh faces. A theme running throughout the thesis is not just an examination of the intrinsic viewpoint, but also the relationship between the intrinsic and the extrinsic approaches.

One of the driving motivations for this investigation into self-Delaunay meshes is the emerging trend in discrete differential geometry towards constructs whose definitions rely on an intrinsic Delaunay triangulation [BS07, DHLM05]. In particular, Wardetzky et al. [WMKG07] identified four properties that would be desirable to have in a discrete Laplacian operator, and then went on to demonstrate that no discrete Laplacian operator can maintain all four properties on arbitrary meshes. However, on self-Delaunay meshes the limitations imposed by that theorem do not apply; Laplacian operators based on the cotan formula [PP93] enjoy all the properties identified by Wardetzky et al.

Formulations of the cotan operator based upon the circumcentric dual complex [Gli05, DHLM05] have been employed on arbitrary meshes when negative edge weights in the Laplacian operator can be tolerated [VL08]. However other complications can arise in this situation. In Appendix A we explain what circumcentric dual cells are, and show that they can have negative area, and discuss possible implications.

One way to avoid these complications and limitations, and to exploit these theories without a self-Delaunay mesh, is to construct an intrinsic Delaunay triangulation of the vertices of the domain mesh. Algorithms to compute and store these intrinsic Delaunay triangulations have been developed [FSBS06], but the resulting data-structure is a complicated compilation of two distinct triangulations, one "overlaid" on the other. The extrinsic
triangulation of the mesh is stored together with the intrinsic Delaunay triangulation and all the points where the edges of the two triangulations cross. The geometry of the domain is defined by piecewise linear interpolation using the extrinsic mesh triangulation, but the intrinsic Delaunay triangulation is used for interpolating data associated with computations involving discrete differential operators. With a self-Delaunay mesh, this complexity is avoided because the two triangulations coincide.

A central theme in discrete differential geometry is the goal of uncovering discrete counterparts of structures and invariants of the smooth theory. This is in contrast to the approach of numerical approximation by discretizing the smooth theory, in which fundamental invariants are no longer strictly preserved in general. Self-Delaunay meshes fit comfortably into this paradigm. Whereas a self-Delaunay mesh is a discrete Delaunay-based structure that may represent $S$, the rDt and the iDt-mesh are Delaunay-based structures only in their roll as discretizations of $S$; they have no independent claim on a Delaunay property.

There are other Delaunay-based mesh structures which have been proposed in the surface reconstruction literature and which do not depend on a separate reference surface for their definition. A notable example of such is the Gabriel mesh, which demands that every triangle face can be inscribed in an Euclidean ball defined by the triangle circumcentre and circumradius and such that no mesh vertices are contained within the ball. In Chapter 4 we examine these structures and demonstrate that Gabriel meshes are self-Delaunay meshes. However, the extrinsic formulation that characterizes Gabriel meshes is more restrictive than what is required of a self-Delaunay mesh, and in general a closed Gabriel mesh won't exist on a given point set $P$, even if $P$ was densely sampled from a smooth surface.

The obstructions to the existence of Gabriel meshes on $P$ are not obstructions for selfDelaunay meshes. However the question of whether or not a self-Delaunay mesh exists on $P$ remains unresolved. It is this issue which motivates our investigation, in Chapter 5, into intrinsic criteria for surface sampling.

It is natural to investigate sampling for a Delaunay structure in terms of the Voronoi diagram, which is dual to the Delaunay triangulation. A Voronoi diagram of a set of sample points $P$ on $S$ is a decomposition of $S$ into Voronoi cells. Each sample $p \in P$ has associated with it a Voronoi cell $\mathcal{V}(p)$ which consists of those points of $S$ which are at least as close to $p$ as to any other sample. What we mean by "as close" depends on which metric we are using to define the Voronoi diagram and the dual Delaunay triangulation. If the metric is the Euclidean distance between points, then the resulting Voronoi diagram is
called the restricted Voronoi diagram ( rVd ), and it is dual to the rDt. An important result of Edelsbrunner and Shah [ES94] states that if the restricted Voronoi diagram of $P$ on $S$ satisfies the closed ball property, then the dual rDt will be a manifold mesh homeomorphic to $S$. For our purposes the closed ball property means that each Voronoi cell is homeomorphic to a disk, and neighbouring Voronoi cells share a single Voronoi edge.

In fact, Edelsbrunner and Shah's result does not depend on the metric employed to define the Voronoi diagram. If the Voronoi diagram satisfies the closed ball property, then the dual simplicial structure will be homeomorphic to the original surface. We exploit this observation to quantify the sampling conditions needed to produce a compatible iDt-mesh from a set of samples on $S$. This problem had previously been investigated by Leibon and Letscher [LL00, Lei99], but the sampling criteria they used were in terms of intrinsic quantities whose relationship with the extrinsic quantities traditionally employed in the surface reconstruction literature was not well understood. We improve on the surface sampling result of Leibon and Letscher, and we simplify the exposition. Further, we quantify the relationship between the intrinsic and extrinsic quantities that are used to govern the sampling density.

These sampling results supply existence conditions for an iDt-mesh, but they do not resolve the issue for self-Delaunay meshes. In Chapter 6 we develop practical algorithms for constructing self-Delaunay meshes from a given input mesh. There is a price to be paid in converting an arbitrary triangle mesh into a self-Delaunay mesh. Either vertices must be added to the original mesh, or some geometric distortion must be tolerated in the conversion to a self-Delaunay mesh. Our viewpoint is that the mesh is intended as a model for a smooth surface, and as such, it is may not be necessary to preserve the precise geometry of the initial mesh. Instead we may aim to remain faithful to the geometry that is being modelled. However, in general the original surface $S$ is not available to us.

The simplest algorithm is to perform extrinsic edge flips on the mesh, in analogy to the famous algorithm of Lawson [Law77] for the planar case. We show that this edge flipping algorithm is guaranteed to terminate because every time an edge between non-coplanar faces is flipped, the surface area of the mesh is reduced. There are two problems with this simple algorithm that manifest themselves on coarse meshes. First, the geometric distortion incurred from flipping an edge with a sharp dihedral angle will be large. Second, and more fundamental, we may encounter unflippable edges: ones that cannot be flipped because the flip would result in a non-manifold edge.

These two deficiencies with the flip algorithm are addressed by another algorithm for constructing self-Delaunay meshes by edge splitting. By using a judicious choice of the split points we are able to prove that the edge splitting algorithm will terminate in a selfDelaunay mesh. The edge splitting algorithm preserves the geometry of the original mesh. Its weakness lies in the number of vertices that are introduced into the final mesh. Other than demonstrating that the number of introduced vertices is finite, we have not quantified a bound on this number. Experiments indicate that it is too large to be interesting in practice, even if an $\mathcal{O}(n)$ asymptotic bound could be established.

The practical algorithm then is a combination of the edge flipping and the edge splitting algorithms. By choosing to only flip those edges whose dihedral angle is sufficiently flat, a balance can be met between introduced geometric distortion and new vertices.

We also introduce a Delaunay mesh decimation algorithm which, given a high resolution self-Delaunay mesh as input, produces a lower resolution mesh that is also guaranteed to be a self-Delaunay mesh. This work is an adaptation of an algorithm for decimating nonobtuse meshes which was introduced by Li and Zhang [LZ06]. The decimation is performed by edge collapse and considering a linearized allowable region in which the resulting vertex must lie in order to maintain the Delaunay property.

The notion of a Delaunay extrinsic edge flip has proven useful not just for constructing self-Delaunay meshes, but also for analyzing their local properties. It is the essential tool employed in comparing Gabriel meshes and self-Delaunay meshes in Chapter 4. The problem of unflippable edges is intimately related to the problem of identifying sampling conditions which ensure the existence of a self-Delaunay mesh. In Chapter 7 we undertake further investigations into the local properties of self-Delaunay meshes, again approaching primarily from the perspective of edge flipping. We show that if a mesh satisfies a mild smoothness constraint, there will be no unflippable edges that are not locally Delaunay.

With this result, the problem of finding sampling conditions which will ensure that a self-Delaunay mesh exists on $P$ can be expressed as a problem of finding criteria which will ensure that the flip algorithm maintains a sufficiently smooth mesh. We make investigations in this direction which broaden our understanding of self-Delaunay meshes, but the central existence problem remains unresolved. In the concluding Chapter 8 we discuss this and other open problems generated by this thesis.

## Chapter 2

## Background

In this chapter we present the established theoretical foundations upon which the rest of the thesis is built. In Section 2.1 we review Voronoi diagrams and Delaunay triangulations as they pertain to Euclidean geometry, outlining their essential properties and applications.

Whereas in the Euclidean domain, the terms triangulation and triangle mesh are practically interchangeable, when we move to the non-Euclidean geometry of smooth surfaces it is prudent to recognize a subtle distinction between these concepts. Such terminology and related definitions are established in Section 2.2.

Then we move on in Section 2.3 to the application of Delaunay concepts to surfaces in $\mathbb{R}^{3}$. There are several ways to extend the basic concepts that are so well developed in the Euclidean domain. Our focus is on mesh structures that have been previously employed in surface reconstruction or approximation theory. In particular, we review the restricted Delaunay triangulation and a closely related structure which we call the intrinsic Delaunay triangulation mesh. Self-Delaunay meshes themselves will be introduced in the next chapter, where they will be compared with the structures we review here.

Finally in Section 2.4 we review the theory of surface approximation with a focus on how it pertains to the Delaunay-based structures. Most of the theory that has been developed in this regard has arisen in the context of surface meshing and reconstruction, and we survey some of the notable works that have contributed to this development.

### 2.1 The Euclidean setting

Voronoi diagrams and Delaunay triangulations are among the most fundamental structures in the computational geometry toolbox. The Voronoi diagram gets its name from the Russian mathematician Georges Voronoi [Vor07, Vor08], although it has appeared under many different names as it has been independently discovered in many different fields. In fact there is evidence that René Descartes was using Voronoi diagrams early in the seventeenth century [Ede01]. The Delaunay triangulation gets its name from another Russian mathematician, Boris Delaunay [Del34], who produced the first systematic modern development of the theory of this structure that is dual to the Voronoi diagram.

These structures are discussed in any computational geometry textbook, see [dBvKOS98] or [Ede01] for example. Comprehensive surveys of the uses, properties and variations of Voronoi diagrams in particular can be found in [AK00] and [OBSC00].

In this section we review the fundamental properties and uses of the Voronoi diagram and the Delaunay triangulation as well as the relationship between them. For the most part, we will confine our attention to Euclidean domains of two or three dimensions, and we focus particularly on properties that will have relevance when we discuss these structures in the context of surface representation.

### 2.1.1 Voronoi diagrams



Figure 2.1: A Voronoi diagram.

The Voronoi diagram is easy to describe and, via a duality relationship, it facilitates the description of the Delaunay triangulation. Given a set $P$ of $n$ points in $\mathbb{R}^{d}$, the Voronoi diagram partitions $\mathbb{R}^{d}$ into $n$ cells: one cell is associated with each point in $P$. For $p \in P$, we denote the associated Voronoi cell by $V(p)$. The extent of $V(p)$ is simply the entire region of $\mathbb{R}^{d}$ whose distance to $P$ is realized by the distance to $p$. That is, the set of points that is at least as close to $p$ as it is to any other $q \in P$. Formally, we have:

Definition 2.1 (Voronoi diagram) The Voronoi cell of $p \in P$ is defined by

$$
\begin{equation*}
V(p)=\left\{x \in \mathbb{R}^{d} \mid d_{\mathbb{R}^{d}}(p, x) \leq d_{\mathbb{R}^{d}}(q, x), \forall q \in P\right\}, \tag{2.1}
\end{equation*}
$$

where $d_{\mathbb{R}^{d}}(p, q)$ denotes the Euclidean distance between $p$ and $q$ in $\mathbb{R}^{d}$. The set of Voronoi cells forms a covering of $\mathbb{R}^{d}$ called the Voronoi diagram of $P$.

The Voronoi diagram gives a very natural definition of the neighbours of a point $p \in P$ : $q$ and $p$ are Voronoi neighbours if their Voronoi cells have a nonempty intersection. This is the relationship that will be exploited to define the Delaunay triangulation in Section 2.1.2.

The Voronoi cells are convex polygons in $\mathbb{R}^{2}$, and in higher dimensions they are convex polytopes. Indeed, $V(p)$ can be constructed as the intersection of the $n-1$ half spaces each of which contains $p$ and is bounded by the orthogonal bisector of $[p, q]$ for some $q \in P$.

Note that we have defined the Voronoi cells to be closed; neighbouring Voronoi cells have a nonempty intersection. The intersection of $d+1$ or more Voronoi cells is either empty or a single point, called a Voronoi vertex. A Voronoi vertex $v$ is equidistant from the elements of $P$ whose Voronoi cells define it. Thus if $v=\cap_{i=0}^{d} V\left(p_{i}\right)$, then the $p_{i}$ all lie on a common hypersphere centred at $v$. For a random set of points $P \subset \mathbb{R}^{d}$, the chances of more than $d+1$ points lying on a common hypersphere is vanishingly small [Ede01]. The set $P$ is said to be in general position if the intersection of more than $d+1$ Voronoi cells is always empty. We will follow the common practice and assume that $P$ is in general position, and in fact this is not a big constraint to impose. If $P$ is not in general position, an arbitrarily small perturbation of the positions of its violating elements is sufficient to bring it into general position. In practice such perturbations are simulated symbolically within an implementation [EM90].

For points in general position, the intersection of $k$ Voronoi cells will be either empty, or a $(d+1)-k$ dimensional polytope called a Voronoi facet. A Voronoi edge arises in the specific case of a non-empty intersection when $k=d$. When $d=3$, a two dimensional Voronoi facet is called a Voronoi face.

In the computational geometry literature, the elements of $P$ are often referred to as sites. For our purposes, we prefer to consider the elements of $P$ to be samples. In the context of geometry processing, the elements of $P$ will be the discrete sample points that represent the geometry (i.e., positional information) of the underlying continuous surface. More generally, a sample set may provide a discrete representation of any signal (or function) defined on the domain.

Suppose then that to each $p_{i} \in P$ there is associated a value $f_{i} \in \mathbb{R}$ and we interpret the $f_{i}$ as the value of some unknown continuous function $f: \mathbb{R}^{d} \rightarrow \mathbb{R}$. The Voronoi diagram
naturally provides an interpolation scheme that allows us to define a function $\tilde{f}: \mathbb{R}^{d} \rightarrow \mathbb{R}$ that coincides with $f$ on $P$. For $x \in V\left(p_{i}\right)$, we simply define $\tilde{f}(x)$ to be $f_{i}$. Although crude and primitive, this interpolation scheme, called nearest neighbour interpolation, is used in some applications and it provides the basis for more sophisticated interpolation and approximation schemes.

A more sophisticated interpolation scheme based on the Voronoi diagram was introduced by Sibson [Sib81]. In this scheme, referred to as natural neighbour interpolation, the value of $\tilde{f}(x)$ is determined by considering the Voronoi cell $V(x)$ of $P \cup\{x\}$. The value of $f$ at $x$ is then given as a weighted sum of the values of $f$ at the sample points $p$, with the weights given by the portion of the original Voronoi cell of $p$ that is occupied by the new Voronoi cell $V(x)$. Thus

$$
\tilde{f}(x)=\sum_{p \in P} \frac{|V(x ;\{x\} \cup P) \cap V(p ; P)|}{|V(p ; P)|} f(p),
$$

where $V(q ; Q)$ denotes the Voronoi cell of $q$ in the Voronoi diagram of $Q$, and $|V|$ represents the size (volume) of set $V$. Sibson demonstrated that this interpolation scheme is smooth everywhere except at the sample points.

Nearest neighbour interpolation is simple, but it is not even continuous. Natural neighbour interpolation is smooth almost everywhere, but it lacks simplicity. A happy medium is obtained with piecewise linear interpolation between the sample points. This involves constructing a triangulation of the sample points and indeed the Voronoi diagram imposes a natural triangulation on the sample set. This triangulation, the Delaunay triangulation, embodies the paradigm upon which this thesis is based.

### 2.1.2 Delaunay triangulations

## Triangulations

A common approach to discretizing a domain is to use a triangulation. A set of samples $P$ is created and these samples are connected by linear elements (line segments, triangles, tetrahedra, etc.) to form a covering of the domain that provides a convenient framework for interpolation and numerical computations. A simplicial complex is a well established structure that allows us to define a triangulation in a unified way for all dimensions. The definition we present here is standard to most textbooks on algebraic topology. See Munkres [Mun84] for example.


Figure 2.2: A simplicial complex is a collection of simplices that satisfies specific intersection rules.

The elements that are the building blocks of a triangulation are called simplices. A set of points $X=\left\{p_{0}, p_{1}, \ldots, p_{m}\right\} \subset \mathbb{R}^{d}$ is affinely independent if the vectors $\left\{p_{i}-p_{0}\right\}_{i \in[1, \ldots, m]}$ are linearly independent. Given such an affinely independent set, the $m$-simplex, $\sigma$, spanned by $p_{0}, \ldots, p_{m}$ is the set of points $x \in \mathbb{R}^{d}$ such that

$$
x=\sum_{i=0}^{m} t_{i} p_{i} \quad \text { with } \sum_{i=0}^{m} t_{i}=1 \text { and } t_{i} \geq 0 \forall i .
$$

In other words, $\sigma$ is the convex hull of the $m+1$ affinely independent points. Generically, when the the dimension need not be specified, we refer to $\sigma$ simply as a simplex. Observe that an $m$-simplex is an $m$-dimensional object. Any subset of the $Y \subset X$, with $k=|Y|-1$, defines an $k$-simplex $\tau \subset \sigma$ that is called an $k$-face of $\sigma$. The 0 -faces of $\sigma$ are the $p_{i}$ themselves and they are called vertices. The 1 -faces are line segments connecting the vertices and they are referred to as edges. Simplices of dimension 2 and 3 are triangles and tetrahedra respectively.

An object $\sigma$ may be similarly defined even if $X$ is not affinely independent. In this case $\sigma$ will not be a truly $m$-dimensional object, and we say that it is a degenerate simplex.

A simplicial complex $\mathcal{K}$ is a collection of simplices such that

1. If $\sigma \subset \mathcal{K}$, then every face $\tau \subset \sigma$ also belongs to $\mathcal{K}$
2. The intersection of any two simplices is a face of each of them.

For the last point, note that the empty set is a face of all simplices. The dimension of $\mathcal{K}$ is the largest dimension of its simplices. Figure 2.2 gives examples of a simplicial complex and a collection of simplices that is not a simplicial complex.

The carrier of $\mathcal{K}$ is the underlying geometric space spanned by the simplices of $\mathcal{K}$, and is denoted $K$. A triangulation of a set of points $P \subset \mathbb{R}^{d}$ is a simplicial complex $\mathcal{K}$ whose vertices are exactly $P$ and such that $K$ forms the convex hull of $P$.

To define a triangulation of a domain $D \subset \mathbb{R}^{d}$ we need to decompose $D$ into a collection of simplices. However, such a decomposition can only be possible if $D$ itself is a polytope, i.e., it has a piecewise linear boundary. If this is not the case, then $D$ must be approximated by some polytope $\tilde{D}$ which can then be triangulated: a triangulation of $\tilde{D}$ is a simplicial complex $\mathcal{K}$ whose vertices include the vertices of $\partial \tilde{D}$, the boundary of $\tilde{D}$, and such that $K$ coincides with $\tilde{D}$.

## Introducing Delaunay triangulations

For $P$ in general position, the Delaunay triangulation of $P \in \mathbb{R}^{d}$ is the dual of the Voronoi diagram of $P \in \mathbb{R}^{d}$. In the planar setting, the duality relationship is as follows: To each Voronoi vertex $c$ we associate a Delaunay triangle, $t$ whose vertices are the three samples which define $c$. An edge $e=[p, q]$ of $t$ is dual to the Voronoi edge $V(p) \cap V(q)$. The vertices of the Delaunay triangulation are the sample points, and they are dual to the corresponding Voronoi cells in the Voronoi diagram.

If $c$ is a Voronoi vertex defined by samples $p, q, s \in P$, then those samples lie on a circle, $C$, centred at $c$ and no samples lie in the interior of the disk bounded by $C$. The circle $C$ is the circumcircle of the Delaunay triangle $t=p q s$. If $P$ is in general position, then each edge of $t$ can be contained in a disk that has the endpoints of the edge as the only two samples on its boundary and that has no samples in its interior. This empty circumdisk property is the standard way to formally define Delaunay triangulations.

Definition 2.2 Let $P \subset \mathbb{R}^{d}$ be a finite set of points. A Delaunay triangulation of $P$ is a triangulation, denoted $\mathcal{D}^{d}(P)$, such that for each $\sigma \in \mathcal{D}^{d}(P)$ there is an open $d$-ball that has the vertices of $\sigma$ on its boundary and which contains no elements of $P$.

If $P$ is in general position, the Delaunay triangulation is known to be unique. If $P$ is not
in general position, the Delaunay triangulation is not uniquely defined. For example, in the plane if four points lie on the boundary of a circle, there are two ways to triangulate their convex hull: both fit the definition of a Delaunay triangulation. Points that are not in general position present a technical annoyance, but as discussed in Section 2.1.1, demanding that $P$ be in general position is not a significant constraint, and we will assume that $P$ is in general position, unless specified otherwise.

## Properties of Delaunay triangulations

In the planar setting, the Delaunay triangulation owes its popularity in part to its ease of definition and construction, but also to the many properties it enjoys. In the plane, the Delaunay triangulation can be constructed with $\mathcal{O}(n \log n)$ operations; in the general case of $d$ dimensions $\mathcal{O}\left(n \log n+n^{\lceil d / 2\rceil}\right)$ operations are required [dBvKOS98].

An early, yet still important algorithm for producing a planar
 Delaunay triangulation was introduced by Lawson [Law72, Law77]. This algorithm takes an arbitrary triangulation of $P$ as input, and produces a Delaunay triangulation by edge flipping. An edge flip replaces an edge $e=[a, b]$ that is adjacent to triangles $[a, b, c]$ and $[b, a, d]$ with the edge $e^{\prime}=[c, d]$ that is the other diagonal of the quadrilateral $[a, d, b, c]$. We refer to $e^{\prime}$ as the opposing edge to $e$. The quadrilateral $[a, d, b, c]$ must be convex for the triangulation to remain valid. We refer to this quadrilateral as the flip-quad associated with $e$.

We say that $e^{\prime}=[c, d]$ is locally Delaunay if $a$ is not contained in the circumcircle of $[b, c, d]$, or equivalently, if $b$ is not contained in the circumcircle of $[c, a, d]$. In his seminal paper, Delaunay [Del34] demonstrated that a triangulation is a Delaunay triangulation if and only if every edge is locally Delaunay. This result applies to Delaunay triangulations of all dimensions, where the locally Delaunay property is defined in terms of circumscribing hyperspheres in the obvious way. A modern exposition of this result using the same proof technique may be found in [Ede01]. Another proof can be found in [dBvKOS98], or [Law77], for example.

There is a convenient characterization of a locally Delaunay edge $e^{\prime}$ : the sum of the angles at the opposing vertices, $a$, and $b$, does not exceed $\pi$ [ BSO 07 ]. With this observation it is easy to show that if edge $e$ is not locally Delaunay (nlD), then $[a, d, b, c]$ is a convex quadrilateral and the opposing edge, $e^{\prime}$, will be locally Delaunay.

Note that an edge that is locally Delaunay in a general triangulation will not necessarily belong to the Delaunay triangulation. The property of being a locally Delaunay edge is not as strong as the property of being a Delaunay edge. Indeed the number of nlD edges in a triangulation is not necessarily a good indicator of how many edge flips are required to achieve a Delaunay triangulation. It is possible to construct a configuration in which there is only one nlD edge, but none of the triangles are Delaunay (see


The red edge is the only nlD edge, but none of the triangles are Delaunay. figure) [GR04].

Lawson's algorithm consists of repeatedly selecting an nlD edge and flipping it. Such a flip is called a Delaunay edge flip. If a nondecreasing sequence $n_{\mathcal{T}}$ is constructed consisting of all the face angles in the triangulation $\mathcal{T}$, then it is shown that a Delaunay flip produces a triangulation $\mathcal{T}^{\prime}$ in which the associated sequence $n_{\mathcal{T}^{\prime}}$ is lexicographically larger than $n_{\mathcal{T}}$ [Law77, Ede01]. Thus the edge flipping algorithm terminates with the triangulation that has the lexicographically maximal sequence of angles. In particular, the resulting triangulation maximizes the minimum angle, and since every edge is locally Delaunay, it is the Delaunay triangulation.

In fact, Lawson first developed his algorithm entirely based upon the criterion of maximizing the minimum angle [Law72]. The connection with the Delaunay triangulation was recognized subsequently [Law77, Sib78].

It turns out that there are several properties, in addition to face angles, that monotonically decrease (or increase) with Delaunay flips. Any such property is optimized (minimized or maximized) by the planar Delaunay triangulation and can be used to demonstrate termination of the edge swapping algorithm. If the property is something that is defined on each triangle, then the optimisation can be expressed in terms of a lexicographic ordering, as described for the triangle face angles.

Consider again the case where each sample $x_{i}$ has an associated value $f_{i}$ that is considered to be the value of some continuous function defined over the domain. Let $f_{\mathcal{T}}$ be the piecewise linear function that interpolates the data over triangulation $\mathcal{T}$. Rippa [Rip90] showed that, regardless of the data $\left\{f_{i}\right\}$, the planar Delaunay triangulation minimizes the Dirichlet energy of $f_{\mathcal{T}}$ :

$$
E\left(f_{\mathcal{T}}\right)=\int_{D}\left\|\nabla f_{\mathcal{T}}\right\|^{2} d a=\sum_{t \in \mathcal{T}}\left\|\left.\nabla f_{\mathcal{T}}\right|_{t}\right\|^{2} a_{t}
$$

where $a_{t}$ is the area of triangle $t$. In other words, the graph of $f_{\mathcal{T}}$ has "minimal roughness" when $\mathcal{T}$ is a Delaunay triangulation.

Another property that is minimized by the planar Delaunay triangulation is the harmonic index [Mus97]. The harmonic index of a triangle, $t$, with sides of length $a, b$ and $c$, is given by $h(t)=\frac{a^{2}+b^{2}+c^{2}}{a_{t}}$, and the harmonic index of a triangulation, $\mathcal{T}$, is given by the sum of the harmonic indices of all the triangles. The harmonic index of $t$ can also be expressed as four times the sum of the cotangents of the angles of $t([\mathrm{BS} 07])$. Note that the harmonic index of a triangle $t$ is minimized if $t$ is equilateral.

The planar Delaunay triangulation also minimizes the maximum circumradius of the triangles [DS89]. As will be discussed in Section 2.4, this is a property that is interesting in the context of surface approximation.

All these properties mentioned have been demonstrated for the planar Delaunay triangulation by exploiting the edge flipping algorithm. However the edge flipping algorithm does not easily extend to higher dimensions, and those extensions that exist are known to not work if started from an arbitrary initial triangulation [ES92]. Thus for most of these properties, it is not known whether or not they extend to a property that is optimized by the higher dimensional Delaunay triangulation.

In some cases, a property that is optimized by the planar Delaunay triangulation is known to be not optimized by higher dimensional Delaunay triangulations. An example is the mean inradius of the triangles. This is maximized by the Delaunay triangulation in 2D, but not the higher dimensions [Lam94].

Using a completely different technique of proof, Rajan [Raj94, Raj91] showed that for all dimensions the Delaunay triangulation minimizes the radius of the maximum smallest enclosing circle. Note that, for obtuse triangles in 2D, for example, this radius is smaller than the circumradius. Rajan's proof does not exploit the flip algorithm, and in fact the Delaunay triangulation does not minimize the smallest enclosing circles in a lexicographic sense.

It is worth mentioning a couple of properties that are not gen-


The nlD edge $[a, b]$ is
shorter than the opposing Delaunay
edge $[c, d]$. erally possessed by the Delaunay triangulation. Contrary to the erroneous assertion in [Lei99][Lemma 24], the Delaunay triangulation does not minimize the total edge length. It is not hard to construct a Delaunay edge flip that increases the
length of an edge (see figure). The problem of constructing a triangulation that minimizes the total edge length is called the minimal weight triangulation problem, and it has recently been shown to be NP-hard [MR06].

The Delaunay triangulation also does not minimize the maximum angle. Although triangulations which possess this property can be desirable for scientific computing applications, computing such triangulations appears more difficult than computing Delaunay triangulations [Ede01][p.12].

## Meshing and Delaunay refinement

The process of producing a triangle mesh to represent a given do-


A cut-away view of a meshed 3D Euclidean
domain main is called meshing. In Section 2.2 we will focus on the problem of meshing the surface of a three dimensional object, however the problem of meshing a Euclidean domain is a huge topic in its own right, and remains an area of active research. Many of the tools and insights discovered in this latter context are beginning to find their way into the field of surface meshing.

A Euclidean domain $D$ is a subset $\mathbb{R}^{d}$ that is a $d$-manifold with boundary. In other words, $D \subset \mathbb{R}^{d}$ is a Euclidean domain if each $x \in D$ has an open neighbourhood in $D$ that is homeomorphic to $\mathbb{R}^{d}$, or to the closed half-space of dimension $d$. The boundary of $D$ is a $(d-1)$-manifold. In this context, the Euclidean space $\mathbb{R}^{d}$ is called the ambient space. So for example, the head show at left ${ }^{1}$ is a Euclidean domain, $D$. The surface of the head is the boundary of $D$ and it is not a Euclidean domain.

The problem of meshing Euclidean domains arises in scientific computing; a PDE is to be solved numerically over the given domain. In this case it is important that the elements of the mesh, the triangle faces if it is a 2D domain, or the tetrahedra if it is a 3D domain, are of good quality. Usually this means that the simplices must be close to regular (maximally symmetric).

The geometry of the problem comes from the boundary of the domain. A mesh must be produced that uses as few vertices as possible given constraints on element size and quality.

[^0]

Figure 2.3: A meshing of Lake Superior, a Euclidean domain. The sample density decreases away from the boundary.

In general there will be smaller elements (more samples) near the boundary of the domain than near the centre (see Figure 2.3) ${ }^{2}$

The dominant paradigm for meshing Euclidean domains is Delaunay refinement. Delaunay refinement is a method for improving the element quality of a given Delaunay mesh by strategically inserting new vertices. Focusing on the 2D case, element quality may be measured by the aspect ratio of the triangles: the ratio between the circumradius and the length of the shortest edge, for example. A Delaunay refinement algorithm will choose the element with the worst (largest) aspect ratio and insert a new sample at its circumcentre. The new sample is then incorporated into the existing mesh and edge flips are performed to maintain the Delaunay property.

A good survey of Delaunay refinement algorithms can be found in Shewchuk's dissertation [She97]. The essential idea of Delaunay refinement is attributed to Rupert [Rup93, Rup95], although a Delaunay refinement algorithm for meshing curved surfaces was independently proposed by Chew about the same time [Che93]. There are many variations

[^1]on the Delaunay refinement paradigm, a notable one being the longest side bisection algorithm of Rivara [RI97], where new vertices are inserted on existing edges, rather than at circumcentres.

In some meshing applications, geometry may be imposed by constraints additional to the boundary of the domain. A constrained Delaunay triangulation of a collection of points and line segments is a triangle mesh that has the points as vertices and the segments as edges and is in some sense as close to being Delaunay as possible, given the constraints on the fixed edges. Chew [Che89] described how to compute a constrained Delaunay triangulation from a given arrangement.

A related concept is a conforming Delaunay triangulation. In this case, we demand that the final triangulation be Delaunay, however the initial constraint line segments can be "split" by the insertion of new vertices. Edelsbrunner and Tan [ET93] demonstrated an algorithm which, given an initial configuration of line segments and vertices, will compute a conforming Delaunay triangulation with an asymptotically optimal (minimal) number of inserted vertices.

In some meshing applications, the sample points are not static, but change location slightly with time. Guibas et al. [GR04] analyzed various algorithms for maintaining Delaunay triangulations under small perturbations of the vertices. The paper makes a number of interesting observations about the fragility of the Delaunay triangulation; a small deformation of the point set may demand a large change in the Delaunay triangulation.

Finally, it should be mentioned that recent work in meshing Euclidean domains has employed alternate metrics to the Euclidean metric. In particular, "geodesic distances" (or "inner distances") measured by lower bounds of lengths of paths within the domain, have proven useful [GGOW08]. In this context the term "Euclidean domain" is arguably no longer appropriate.

### 2.1.3 Generalizations and related structures

The Voronoi-Delaunay paradigm has spawned many variations and related structures. The basic Delaunay structures for surface representation discussed in Section 2.3 represent extensions of this paradigm to non-Euclidean domains. For now, we remain within the context of Euclidean domains and review a few more structures that are related to the VoronoiDelaunay structures introduced above.

## Power diagram

One natural generalization of the Delaunay triangulation is the weighted Delaunay triangulation. It is sometimes referred to as a regular triangulation, but this latter terminology has been applied also to quite different concepts (c.f. [FSBS06]) and so we will avoid it.

In a weighted Delaunay triangulation, each vertex $p$ is assigned a weight $w_{p}$. The weighted Delaunay triangulation can then be described as the dual of the power diagram of $P_{w}=\left\{\left(p, w_{p}\right)\right\}$. The power diagram is analogous to the Voronoi diagram, but the distance associated with each sample is dependent on its weight. The power distance from $\left(p, w_{p}\right)$ to $x \in D$ is given by $\pi_{p}(x)=d_{\mathbb{R}^{3}}(x, p)^{2}-w_{p}^{2}$. Then the power cell associated with $\left(p, w_{p}\right)$ is given by

$$
V_{w}(p)=\left\{x \in D \mid \pi_{p}(x) \leq \pi_{q}(x) \quad \forall q \in P_{w}\right\} .
$$

It can be shown that these power cells are convex polyhedra and the power diagram is the collection of all the power cells. The weighted Delaunay triangulation is the dual of the power diagram in the same way that the usual Delaunay triangulation is the dual of the Voronoi diagram. In fact, if all the weights are the same, the power diagram is exactly the Voronoi diagram described in Definition 2.1, and the corresponding weighted Delaunay triangulation is just the ordinary Delaunay triangulation.

One geometric interpretation of the power diagram is obtained by considering $w_{p}$ to be the radius of a sphere centred at $p$. The power distance to $x$ is then the square of the length of the line segment from $x$ to a tangent point on the sphere. More details can be found in [Dey07] or [Ede01].

An interesting power diagram can be constructed as follows: Let $P \subset \mathbb{R}^{d}$ and let $V$ be the set of Voronoi vertices in the Voronoi diagram of $P$. Assign to each Voronoi vertex, $v$, a weight, $w_{v}$, equal to the radius of the circumsphere of the Delaunay simplex dual to $v$. Then the power diagram of $V_{w}=\left\{\left(v, w_{v}\right)\right\}$ is exactly the Delaunay triangulation of $P$ [Dey07][§10.4].

## Alternate metrics

The power diagram represents a distortion of the Voronoi diagram which maintains the convex polyhedron property of the cells. Another natural generalization of the Voronoi diagram of Definition 2.1, is obtained by replacing the Euclidean metric $d_{\mathbb{R}^{3}}$ with an alternate metric. In this case convexity of the Voronoi cells is no longer assured. In fact for general metrics
many desirable properties of the Voronoi cells can be lost, and even defining the Voronoi diagram can be a challenge. For example, the Manhattan $\left(L_{1}\right)$ metric in the plane can yield Voronoi cell bisectors that are regions and not curves [KW88]. Klein and Wood [KW88] developed a set of axioms that, when satisfied by a metric, assure that the Voronoi diagram in the plane is well behaved. Aurenhammer and Klein [AK00] refer to this class of metrics as nice metrics. In particular, nice metrics yield connected Voronoi cells with boundaries composed of a finite set of Voronoi cell bisectors that are simple curves.

An nice example of Delaunay meshes constructed using an alternate metric is in the work of Labelle and Shewchuk [LS03], in which they studied anisotropic metrics. As they remark, the setting they consider is (at least locally) equivalent to the case involving the intrinsic metric of a surface, as will be discussed in Section 2.3.3. However, for computational reasons they consider an approximation to the Voronoi diagram of the Riemannian metric, and in so doing they lose some of its nice properties. In particular, the Voronoi cells may not be connected.

## Gabriel complexes

Another structure built on the Delaunay theme that will be of interest to us is the Gabriel complex. This is actually a substructure of the usual Delaunay triangulation, and we will study it in some detail in Chapter 4. Recall that the Delaunay triangulation consists of those simplices that can be contained in a ball whose interior is empty of samples. The $k$ Gabriel complex is the simplicial complex which consists of those $k$-simplices for which the smallest circumscribing ball (i.e., the smallest ball that has all the vertices of the simplex on its boundary) is empty, together with the faces of these simplices. Thus in $\mathbb{R}^{d}$, the $d$-Gabriel complex is the entire Delaunay triangulation. The 1-Gabriel complex is called the Gabriel graph (see Figure 2.4).

A description of the Gabriel graph was published in 1969 [GS69], and it is now a well established and much used data structure. However, the higher dimensional analogue has a much shorter history. The 2 -Gabriel complex in $\mathbb{R}^{3}$ was introduced by Petitjean and Boyer [PB01].


Figure 2.4: The Gabriel graph consists of those Delaunay edges that intersect their dual Voronoi edges. These edges are shown here in magenta, whereas the Delaunay edges that are not Gabriel are cyan. The Voronoi edges are drawn with dashed red lines.

### 2.2 Triangulations and triangle meshes for surfaces

We now turn our attention to non-Euclidean domains. In particular, our interest is in the 2D surfaces that bound objects in 3D space, however we will also need to deal with surfaces that cannot be described in this way, because they have self-intersections for example, and this will lead us to more abstract concepts. Our focus is on smooth surfaces, and triangle meshes intended to model such surfaces. The purpose of this section is to establish our terminology, and to build a careful definition of a manifold triangle mesh. In the process we distinguish several closely related concepts.

### 2.2.1 Embeddings and Immersions

Our investigations involve surfaces of various kinds. The question of whether two objects (surfaces for example), $\mathcal{F}$, and $\mathcal{F}^{\prime}$, are the same may have a different answer, depending on the amount of structure we are considering. For example, $\mathcal{F}$ and $\mathcal{F}^{\prime}$ may be topologically equivalent, but they may each possess different metric structures, so that they are not geometrically equivalent.

A mapping $h: \mathcal{F} \rightarrow \mathcal{F}^{\prime}$ is injective (one to one) if no two points in $\mathcal{F}$ get mapped to the same point in $\mathcal{F}^{\prime}$. It is surjective if every point in $\mathcal{F}^{\prime}$ is the image under $h$ of some point in $\mathcal{F}$. If $h$ is both injective and surjective, then it is bijective.

Our objects, $\mathcal{F}$ and $\mathcal{F}^{\prime}$, are considered equivalent if there exists a bijective mapping between them which preserves the structures under consideration. Generically, such a mapping is called an isomorphism, and we say that $\mathcal{F}$ and $\mathcal{F}^{\prime}$ are isomorphic. Specifically, if we are considering only the topological structure of $\mathcal{F}$ and $\mathcal{F}^{\prime}$, an isomorphism is called a homeomorphism. A bijective mapping $h$ is a homeomorphism if it is continuous and it has a continuous inverse, i.e., $h^{-1}: \mathcal{F}^{\prime} \rightarrow \mathcal{F}$ must also be continuous. In this context, $\mathcal{F}$ and $\mathcal{F}^{\prime}$ are said to be homeomorphic: they are topologically equivalent.

If $\mathcal{F}$ and $\mathcal{F}^{\prime}$ are differentiable manifolds, then an isomorphism between them is called a diffeomorphism. A diffeomorphism is a differentiable bijective mapping whose inverse is also differentiable. If we are considering in addition a metric structure on $\mathcal{F}$ and $\mathcal{F}^{\prime}$, then an isomorphism must also preserve distances between corresponding points. Such a mapping is called an isometry.

Often we view our object, $\mathcal{F}$, as an entity in $\mathbb{R}^{3}$, the ambient space. An embedding $\phi: \mathcal{F} \rightarrow \mathbb{R}^{d}$ is a mapping in which the domain is isomorphic to the image, i.e., $\phi$ is an isomorphism between $\mathcal{F}$ and $F=\phi(\mathcal{F}) \subset \mathbb{R}^{d}$. If $\mathcal{F}$ is embedded in $\mathbb{R}^{d}$, then the distinction between $\mathcal{F}$ and $F$ is mathematically unimportant.

The mapping $\phi$ is an immersion if it is locally an embedding. This means that each point $x \in \mathcal{F}$ admits an open neighbourhood $V$ such that $\left.\phi\right|_{V}$ is an embedding, where $\left.\phi\right|_{V}$ is the restriction of $\phi$ to $V$. For the particular case of a smooth mapping $\phi$, the inverse function theorem (see [Boo75][p.42] for example) ensures that a necessary and sufficient condition for $\phi$ to be an immersion is that it's Jacobian has full rank everywhere.

The criterion for an immersion is weaker than is demanded by an embedding. In particular, if $\phi$ is an immersion, it need not be injective. The concept of an immersion allows us to consider surfaces with self intersections, as described in Figure 2.5.

A more detailed exposition of embeddings and immersions can be found in any textbook on advanced calculus. A nice geometry oriented overview can be found in Chapter 2 of [Boo75].


Figure 2.5: An embedding and an immersion. (a) This torus is topologically embedded in $\mathbb{R}^{3}$; it has no self-intersections. (b) Here the torus is not embedded since it has selfintersectons (it is twisted into a figure eight). Topologically, this object is not a torus, however we can view it as a torus immersed in $\mathbb{R}^{3}$.

### 2.2.2 Surfaces

## Smooth surfaces

In this exposition a smooth surface is a smooth $\left(C^{\infty}\right)$, manifold surface $S \subset \mathbb{R}^{3}$. Such surfaces are often called regular surfaces in textbooks on classical differential geometry [dC76], where they are the principle objects of study. Such a textbook should be consulted for details on the definitions and assertions made here.

Unless specified otherwise, we assume that $S$ is the boundary of a Euclidean domain. This means that $S$ is compact, connected, orientable and has no boundary. Thus a smooth surface is a Riemannian 2-manifold $\mathcal{S}$ isometrically embedded in $\mathbb{R}^{3}$. Note that if $\mathcal{S}$ were not orientable, then it would not bound a Euclidean domain, and indeed would not admit an embedding into $\mathbb{R}^{3}$, however it may be possible to find an immersion into $\mathbb{R}^{3}$, as demonstrated


Klein bottle by the Klein bottle shown here.

A curve on $S$ is an immersion $\gamma: I \subset \mathbb{R} \rightarrow S$, and such a curve may also be viewed as a space curve: $I \rightarrow \mathbb{R}^{3}$. All curves will be parameterized by arclength. Thus the tangent
vector to $\gamma$ has unit length: $\left\|\gamma^{\prime}(s)\right\|=1$ for all $s \in I$. The length of $\gamma$ is $\ell(\gamma)$, and we abuse the notation by identifying $\gamma$ with its image. Thus for example, $z \in \gamma$ means $z \in \gamma([0, \ell(\gamma)])$.

The normal curvature at $p$ of a curve $\gamma \subset S$ is the scalar product between $\hat{\boldsymbol{n}}_{S}(p)$, the normal to $S$ at $p$, and $\gamma^{\prime \prime}(s)$, the curvature vector of $\gamma$ at $p=\gamma(s)$. Each point $p \in S$ has associated with it two principle curvatures, $\kappa_{1}(p), \kappa_{2}(p)$ given by the maximum and minimum of the normal curvatures of all curves through $p$. The maximal curvature at $p$ is the maximum of the absolute values of the two principal curvatures and is denoted $\kappa(p)$. The Gaussian curvature is the product of the principle curvatures, and is denoted $G(p)$.

A geodesic on $S$ is a curve $\gamma \subset S$ that, when viewed as a space curve $\gamma \subset \mathbb{R}^{3}$, has its curvature vector parallel to the normal vector on $S$ at all corresponding points where the curvature of $\gamma$ does not vanish. In particular, the curvature of $\gamma$ at $x \in \gamma$ is bounded by the largest magnitude of the principle curvatures of $S$ at $x$.

Compact surfaces are geodesically complete; any two points $p, q \in S$ can be connected by a smooth curve $\gamma \subset S$ of minimal length, and this curve is a geodesic (the Hopf-Rinow theorem). Also, for any $z \in \gamma$, the portion of $\gamma$ between $p$ and $z$ is also a minimal geodesic between $p$ and $z$ [dC92]. However, a minimal geodesic is not necessarily unique. If two points are sufficiently close together we can be assured of a unique minimal geodesic between them [dC92].

The geodesic distance between two points $p, q \in S$, denoted $d_{S}(p, q)$, is the length of the shortest path between them. If $S$ is viewed as a Riemannian manifold, the metric $d_{S}$ arises naturally from its Riemannian metric tensor and as such does not depend on the embedding in $\mathbb{R}^{3}$ for its definition. (Although the embedding can be used to define the Riemannian metric tensor, this tensor remains well defined in the absence of any embedding.) We call this metric the intrinsic metric on $S$ and likewise all properties which depend only on this metric for their definition, are deemed intrinsic. A property which depends on the embedding of $S$ in the ambient space, is an extrinsic property.

## Piecewise flat surfaces

In Section 2.2.4 we will introduce manifold triangle meshes, but these objects are not smooth surfaces. Just as it is sometimes convenient to consider $S$ as an abstract smooth 2-manifold $\mathcal{S}$, divorced from any particular embedding in $\mathbb{R}^{3}$, it is useful to have a corresponding abstraction for meshes.

A piecewise flat surface ( pwf surface), $\mathcal{M}$, is a two dimensional manifold equipped with


Figure 2.6: Cones and pwf surfaces. (a) The lateral surface of a cone is the model for the local geometry of a pwf surface. Any point $x$ that is not the apex of the cone, $p$, admits an open neighbourhood (blue) that is isometric to an open set on the plane. (b) The cone may be developed onto the plane by inserting a cut from the perimiter of the base to the apex. Note however that in the definition of a pwf surface we are not restricting ourselves to convex cones; the perimeter of the developed cone may subtend an angle greater than or equal to $2 \pi$ at $p$. Thus a planar disk is also a cone. If a point on a pwf surface does not admit a neighbourhood isometric to a planar disk, then it will admit a neigbourhood isometric to a neigbourhood of the apex of a cone (red region in (a)). Such a point is a cone point. (c) This definition encompasses the intrinsic geometry of polyhedra. The edges of a polyhedron are extrinsic artifacts; every point in a polyhedral surface admits a neighbourhood isometric to the lateral surface of a cone.
a metric such that each point $x \in \mathcal{M}$ has a neighbourhood isometric to the lateral surface of a cone (see Figure 2.6). A comprehensive study of the geometry of these objects can be found in the treatise by Aleksandrov and Zalgaller [AZ67]. The point $x$ is a cone point if it does not admit a neighbourhood isometric to a planar disk. The definition implies that the cone points are isolated. We will only consider compact pwf surfaces and therefore the number of cone points will necessarily be finite.

A two dimensional manifold simplicial complex is an example of a pwf surface; the vertices of the simplicial complex include all the cone points. It is in fact a properly triangulated pwf surface, as we now explain.

### 2.2.3 Surface triangulations

A triangulation of a two-manifold, $\mathcal{S}$, is a simplicial complex $\mathcal{K}$ together with a homeomorphism $h: K \rightarrow \mathcal{S}$. The vertices, edges, and triangles of the triangulation are the images of the vertices, edges and triangles of $\mathcal{K}$. If a set of samples $P \subset S$ is given, then a triangulation of $P$ on $S$ is a triangulation $h: \mathcal{K} \rightarrow S$ such that the vertices of $\mathcal{K}$ are mapped bijectively onto $P$. It can be shown that the surfaces we consider always admit a triangulation, however it is not true that any set of samples $P \subset S$ is sufficient to be the vertex set of a triangulation. The topology of the surface will demand a minimum number of samples, an issue which we will explore in some detail in Chapter 5.

In the case of a piecewise flat surface, $\mathcal{S}$, we demand that the homeomorphism $h$ be an isometry. In Section 3.1 we will consider a weaker notion of a triangulation of a pwf surface, so we refer to the triangulation as defined here as a proper triangulation of the pwf surface. Thus a properly triangulated pwf surface is a two dimensional manifold simplicial complex, and we may use the terms interchangeably. The vertices of a properly triangulated pwf surface include all of the cone points. The star of a vertex $p$ is the union of the simplices incident to $p$, including their faces.

### 2.2.4 Triangle meshes

A piecewise linear function on a triangulated pwf surface is a function that is linear on each simplex. Such a definition is possible, because a triangulated pwf surface may be isometrically identified with a two dimensional manifold simplical complex (Figure 2.7), and the definition of a linear function on a Euclidean simplex is clear.


Figure 2.7: A triangle in a triangulated pwf surface is isometric to a Euclidean triangle. gether with a piecewise linear mapping $\phi: \mathcal{K}_{M} \rightarrow \mathbb{R}^{3}$, whose restriction to each simplex is an isometric linear embedding. If $\phi$ is an immersion, then $M$ is a manifold triangle mesh, otherwise, it is a singular triangle mesh. If $\phi$ is an embedding, then $M$ is an embedded mesh. We generally identify $M$ with the image of $\phi$.

Since we do not insist that $\phi$ be an embedding, an arbitrary $M$ may have self-intersections, as in Figure 2.5(b), for example. However, for all meshes the underlying simplicial complex, $\mathcal{K}_{M}$, is always a 2 -manifold. In the case of a manifold triangle mesh, the requirement that $\phi$ be an immersion implies that no two triangles in the star of any vertex may intersect in their interiors: the star of any vertex $p \in M$ is homeomorphic to a closed disk (see Figure 2.8). We refer to the star of $p$ as the umbrella at $p$ so as to emphasise the manifold property and we denote it $U(p)$. The one ring of $p$ is the set of edges and vertices that comprise the boundary of $U(p)$.


Figure 2.8: If two triangles in the star of a vertex $p \in M$ intersect, as do $[p, u, w]$ and [ $p, q, v$ ] shown here, then $M$ is a singular triangle mesh. Although our definition of a manifold triangle mesh encompasses surface meshes that self-intersect, such as that shown in Figure 2.5(b), the requirement that $M$ be the image of an immersed pwf surface prohibits local self-intersections such as is shown here. Indeed, every point in a pwf surface admits a neighbourhood homeomorphic to a topological disk and whose image under $\phi$ is also a topological disk. However, in the example shown here, $p$ admits no neighbourhood in $M$ homeomorphic to a disk. We will see that this distinction is important in the context of extrinsic edge flipping: Proposition 7.5 does not apply to singular triangle meshes.

We will focus on manifold triangles meshes without boundary, which means that every vertex lies in the topological interior of the disk defined by its star. In the event that a mesh with boundary is considered, a boundary vertex is one which lies on the boundary of
its star ${ }^{3}$. If $p$ is an interior vertex, we may emphasize this by saying it has a full umbrella. An edge is a boundary edge if it connects two boundary vertices.

Since self intersections may occur, $M$ is in general not a simplicial complex as defined in Section 2.1. However, the combinatorial topology that is inherent in the associated simplicial complex, $\mathcal{K}_{M}$, is preserved. An abstract simplicial complex $\mathcal{A}$ is a collection of finite nonempty sets such that if $C \in \mathcal{A}$, then so is every nonempty subset of $C$. An element $C$ of $\mathcal{A}$ is called a simplex of $\mathcal{A}$. The dimension of a simplex is one less than the number of elements it contains. So for example, the collection of simplices in Figure 2.2(b) has the structure of an abstract simplicial complex, even though it is not a simplicial complex.

Praun et al. [PSS01] ${ }^{4}$ define a triangle mesh as a pair $M=\left(P, \mathcal{A}_{M}\right)$, where $P$ is a set of $n$ point positions in $\mathbb{R}^{3}$ and $\mathcal{A}_{M}$ is an abstract simplicial complex composed of three types of subsets of $\{1, \ldots, n\}$ : vertices $\{i\}$, edges $\{i, j\}$ and faces $\{i, j, k\}$.

A geometric realization of an abstract simplicial complex, $\mathcal{A}$, is a simplicial complex $\mathcal{K}$ together with a bijection $h$ between the vertices of $\mathcal{A}$ and the vertices of $\mathcal{K}$ and such that $C \in \mathcal{A}$ iff the simplex spanned by $h(C)$ is in $\mathcal{K}$. Thus our definition of a triangle mesh agrees with that of Praun et al., but with the added stipulation that the carrier of a geometric realization of $\mathcal{A}_{M}$ be a 2 -manifold.

Generally the geometric realization is a means of describing the global topology associated with $\mathcal{A}$. However, in the case of a triangle mesh at least, we are very much interested in the geometry as well as the topology. Thus the canonical geometric realization of $\mathcal{A}_{M}$ is the properly triangulated pwf surface $\mathcal{K}_{M}$ introduced in the definition above. The carrier of $\mathcal{K}_{M}$, the unadorned pwf surface associated with $M$, is denoted $\mathcal{M}$.

## Qualities of meshes representing a smooth surface

Ideally we wish to produce a manifold triangle mesh $M$ that is a good approximation to $S$. This problem, of producing a mesh to represent a given surface is known as surface meshing and has seen considerable attention in the literature. A closely related problem is to produce a mesh $M$ to represent $S$ when $S$ itself is unknown, except for the discrete set of sample points $P$, which are to be the vertices of $M$. This problem is known as surface reconstruction. To summarize the distinction between the two problems: in surface meshing

[^2]$S$ is given, and the problem is to generate $P$ and $M$, whereas in surface reconstruction, only $P$ is given, and the problem is to produce $M$. In both cases of course it is desired that $M$ be a good representation of $S$. Most of the mesh structures and surface approximation theory that we review in this chapter has arisen from works in meshing or surface reconstruction.

Exactly what is meant by a good approximation will be discussed in detail in Section 2.4. Certainly, a good mesh will be an embedded one which is homeomorphic to $S$. However, we will also be considering a weaker conformity to $S$. We say that $M$ triangulates $S$ if $\mathcal{M}$ is homeomorphic to $S$. In this case, $M$ itself may not be embedded.

Given a vertex $p$ in a manifold triangle mesh, $M$, we arbitrarily choose a consistent orientation for the face normals in the umbrella of $p$. We say that $p$ (or its umbrella) is $\theta$-smooth if for any two triangles $t_{1}, t_{2}$ incident to $p$, the angle between their normal vectors is less than $\theta$. The mesh itself is $\theta$-smooth if every vertex of $M$ is $\theta$-smooth. A smooth mesh is a manifold triangle mesh that is $\theta$-smooth for some $\theta<\pi / 2$. Likewise, an unqualified reference to a smooth vertex, or smooth umbrella implies that it is $\theta$-smooth for some $\theta<\pi / 2$. We emphasize the requirement that the mesh must be non-singular in order to be called a smooth mesh. This constraint is not a consequence of the constraints on the triangle normal vectors.

As we will discuss in Section 2.4, a mesh that is a good approximation to $S$ will be smooth. In Chapter 7 we will also show that smooth meshes are a natural domain for the extrinsic edge flipping algorithm which we will introduce in Section 3.2.1.

### 2.3 Delaunay structures for surfaces

In this section we review methods of employing the Delaunay paradigm to define meshes that represent $S \subset \mathbb{R}^{3}$. We begin in Section 2.3 .1 by extending to surfaces the definition of a Voronoi diagram that we introduced in Section 2.1 in the context of Euclidean space. Extending this definition is trivial, however extending the definition of Delaunay triangulations to surfaces faces complications. We study conditions on the Voronoi diagram which, when met, ensure the existence of a mesh which is naturally dual to the Voronoi diagram and which triangulates $S$.

When the conditions for the existence of such a mesh exist, then as with traditional Delaunay triangulations, the faces and edges may equivalently be described directly by an empty circumdisk criterion, rather than as the dual to the Voronoi diagram. The Voronoi
diagram, and the empty circumdisks are defined by a choice of metric, i.e. a decision on how to measure distances between points on $S$. We examine in particular two distinct structures arising from different choices of metric.

In Section 2.3.2 we discuss the restricted Delaunay triangulation (rDt) which results from the criterion that each mesh triangle face $t$ be contained in a (3D) Euclidean ball centred on $S$ and such that the vertices of $t$ are on its boundary and no sample points of $P$ are in its interior. This is just the empty circumdisk property when $d_{\mathbb{R}^{3}}$ is employed to measure distances on $S$. Then, in Section 2.3.3 we look at the mesh that results when $d_{S}$, the intrinsic metric on $S$ is employed.

Other structures which have a Delaunay flavour, but which don't completely fit the Delaunay paradigm are reviewed in Section 2.3.4.

### 2.3.1 Voronoi diagrams on surfaces

In exact analogy with Definition 2.1, a Voronoi diagram of a sample set $P$ on $S$ is naturally defined without imposing restrictions on $P$.

Definition 2.4 (Voronoi diagram on $S$ ) The intrinsic Voronoi cell of $p \in P$ is defined by

$$
\begin{equation*}
\mathcal{V}(p)=\left\{x \in S \mid d_{S}(p, x) \leq d_{S}(q, x), \forall q \in P\right\} . \tag{2.2}
\end{equation*}
$$

The collection of intrinsic Voronoi cells is called the intrinsic Voronoi diagram (iVd) of $P$ on $S$. Replacing $d_{S}$ with $d_{\mathbb{R}^{3}}$ in Equation (2.2), we obtain the restricted Voronoi cell of $p$ on $S$, which is denoted $\left.V\right|_{S}(p)$, because it is precisely the restriction to $S$ of $V(p)$, the Euclidean Voronoi cell of $p \in P \subset \mathbb{R}^{3}$. The resulting collection of Voronoi cells is called the restricted Voronoi diagram (rVd) of $P$ on $S$.

## The nerve of a covering

A covering of a topological space $X$, is a collection $C$ of subsets of $X$ such that $X=\cup_{A \in C} A$. $C$ is a closed (open) covering if every set in $C$ is closed (open). Thus a Voronoi diagram on $S$ is a closed covering of $S$.

The nerve of a covering, $C$, is an abstract simplicial complex $\mathcal{A}$ whose simplices are defined by subsets of $C$ which have a nonempty mutual intersection. Thus $\mathcal{A}=\{D \subset$ $\left.C \mid \cap_{A \in D} A \neq \emptyset\right\}$. A geometric realization of the nerve of a covering is a simplicial complex


Figure 2.9: A covering and its nerve. (a) A covering, $C$, of the set $F$ (yellow) is the collection of sets $A_{i} \cap F$, where each $A_{i}$ is depicted by a coloured loop. (b) This simplicial complex is a geometric realization of the nerve of $C$.
that is in a sense dual to the covering (see Figure 2.9). For the iVd, this dual complex is called the intrinsic Delaunay complex.

The dual complex to the rVd is the restricted Delaunay complex, and in this case, the geometric realization is canonically given by a subcomplex of $\mathcal{D}^{3}(P)$. This subcomplex is defined by those Delaunay triangles whose dual Voronoi edges intersect $S$, together with those Delaunay edges whose dual Voronoi facets intersect $S$. The vertex bijection defining the geometric realization is then given by $h\left(\left.V\right|_{S}(p)\right)=p$.

We are interested in the conditions under which the carriers of these dual complexes will be homeomorphic to $S$.

## The closed ball property

Edelsbrunner and Shah [ES94] introduced the closed ball property to describe those restricted Voronoi diagrams on $S$ whose restricted Delaunay complex is homeomorphic to $S$. For a surface without boundary, the property expresses three conditions:

1. each Voronoi cell is a closed topological disk (2-ball);
2. the intersection of two Voronoi cells is either empty or a closed topological 1-ball: a single Voronoi edge;


Figure 2.10: Two Voronoi cells on the dolphin's left fin share two Voronoi edges, a violation of the closed ball property.
3. the intersection of three Voronoi cells is either empty or a single point (0-ball): a single Voronoi vertex.

Figure 2.10 shows an example of a violation of the second closed ball property.

General position: The closed ball property implicitly imposes a general position condition on the sample points in that the intersection of more than three Voronoi cells must necessarily be empty (otherwise there would be two Voronoi cells whose intersection was a single point). No sampling criteria based on sample density alone can guarantee that the general position property is satisfied. Edelsbrunner and Shah [ES94] gave arguments to support the assumption of general position (they use the term general intersection property) in the rVd. These arguments don't apply to the iVd, but we supply in Section C. 1 evidence that the assumption of general position is also reasonable in this case, at least when $S$ is a compact surface. Henceforth we will assume that $P$ is in general position for whichever Voronoi diagram on $S$ is under consideration.

Generality: Edelsbrunner and Shah [ES94] showed that the restricted Delaunay complex is homeomorphic to $S$ when the rVd satisfies the closed ball property. Their proof looks at the topological properties of the dual complex resulting from a Voronoi diagram satisfying
this property. An examination of that proof reveals that it does not rely on the specific metric employed to generate the Voronoi diagram. In particular, if the iVd satisfies the closed ball property, then the intrinsic Delaunay complex will be homeomorphic to $S$.

Redundancy of condition 3: We have assumed that $S$ is a single component (it is connected). We further assume that there are at least four distinct samples in $P$. In this case, the third condition of the closed ball property becomes redundant. Indeed, suppose that the first two closed ball conditions are satisfied but the third is violated. Thus we have three Voronoi cells $\mathcal{V}(p), \mathcal{V}(q)$ and $\mathcal{V}(r)$ whose intersection includes two distinct points $a$ and $b$. Then $\mathcal{V}(p) \cap \mathcal{V}(q)$ will be a Voronoi edge, $e$, following condition 2, and $a$ and $b$ must necessarily be its endpoints. Likewise, $\mathcal{V}(p) \cap \mathcal{V}(r)$ will yield another distinct Voronoi edge, $e^{\prime}$, also with $a$ and $b$ as endpoints. Since $\mathcal{V}(p)$ is a topological disk (condition 1 ), $e$ and $e^{\prime}$ make up its entire boundary. It follows that $p$ has no Voronoi neighbours other than $q$ and $r$. Arguing similarly for $q$ and $r$, we conclude that there are only three samples on $S$.

Definition 2.5 (Well formed Voronoi diagram) A (restricted or intrinsic) Voronoi diagram is well formed if it consists of at least four samples and satisfies the first two closed ball conditions.

### 2.3.2 Restricted Delaunay triangulation (rDt)

When the rVd is well formed, the restricted Delaunay complex is a triangle mesh known as the restricted Delaunay triangulation (rDt). Edelsbrunner and Shah's result [ES94] ensures that the rDt is an embedded mesh that triangulates $S$. However, in spite of its name, it is a mesh and only becomes a triangulation of $S$ when a specific homeomorphism onto $S$ is chosen. As we will discuss in Section 2.4, it is typical to seek sampling conditions which ensure that the projection onto the closest point on $S$ provides the triangulation homeomorphism.

The restricted Delaunay triangulation (rDt) has been extensively studied in the geometry processing community. It is the most commonly used Delaunay structure for surface representation. It was first formally defined by Edelsbrunner and Shah [ES94] via the duality with the rVd. However, the rDt appeared in an earlier meshing algorithm by Chew [Che93], but no explicit reference to the rVd was made in that work. Instead, the faces of the mesh were defined directly in terms of empty Euclidean spheres.


Figure 2.11: The restricted Delaunay triangles for a sample set $P \subset S$ are characterized by the property that they admit an empty Euclidean circumball centred on $S$, depicted by the blue circle in this one dimensional example. The centre of this ball, $c$, is a Voronoi vertex in the rVd , and it is also the point at which the dual Voronoi edge in the three dimensional Voronoi diagram of $P$ intersects the surface. In this one dimensional example, it is the red edge labeled $V(p q)$.

Intuitively, the rDt is the mesh that is obtained when Voronoi neighbours in the rVd are connected by mesh edges; the faces of the rDt are Euclidean simplices dual to the Voronoi vertices in the rVd. The rDt can be equivalently characterized as consisting of those triangle faces of $\mathcal{D}^{3}(P)$ for which the dual Voronoi edge intersects $S$. The point of intersection is a Voronoi vertex in the rVd. This latter characterization leads to the property, exploited by Chew [Che93], that the faces of the rDt are contained in empty Euclidean spheres centred on $S$ (see Figure 2.11).

Since the initial works of Chew and Edelsbrunner and Shah, many meshing and surface reconstruction algorithms have been based upon the rDt and an analysis of the rVd. In particular, the first provably correct surface reconstruction algorithms [AB98, ACDL00] exploited the work of Edelsbrunner and Shah. Chew's work has inspired many subsequent algorithms, not just in the meshing of smooth surfaces specifically [BO05], but also in surface sampling [BGO05], in polygonal remeshing [DLR05], and the meshing of a wider class of non-smooth surfaces [BO06].

### 2.3.3 Intrinsic Delaunay triangulation mesh (iDt-mesh)

When the iVd is well formed, then the intrinsic Delaunay complex is homeomorphic to $S$. The obvious piecewise linear mapping that identifies the vertices of the intrinsic Delaunay complex with the corresponding sample points in $P$ defines a triangle mesh which we call the iDt-mesh (short for intrinsic Delaunay triangulation mesh). Vertices are neighbours in the iDt-mesh if they are Voronoi neighbours in the iVd on $S$.

The iDt-mesh triangulates $S$, but unlike the rDt , there is no guarantee that it will be embedded, or even immersed. Such guarantees would require that extrinsic sampling criteria are imposed on the sample set; the relationship between intrinsic and extrinsic sampling criteria are discussed in Chapter 5.

We distinguish the iDt-mesh from the intrinsic Delaunay triangulation (iDt) of $P$ on $S$, which is not a mesh, but a triangulation of $S$ in which the edges are minimal geodesics between the vertices. Such a triangulation is called a geodesic triangulation ${ }^{5}$ of $S$. With this stipulation on the iDt, the closed ball property of the iVd is not sufficient to ensure that the iDt is well defined: We need the further assurance that each pair of neighbouring vertices are connected by a unique minimal geodesic on $S$.

A geodesic disk centred at $c \in S$ with radius $r$ is the set $B_{S}(c ; r)=\left\{x \in S \mid d_{S}(c, x)<r\right\}$. We work primarily with open disks and balls, but if we wish to refer to a closed geodesic disk, it will be denoted $\bar{B}_{S}(c ; r)$, and more generally, $\bar{A}$ denotes the topological closure of set $A$. In general a geodesic disk need not even be a topological disk (it can wrap around on itself), but for any $c \in S$, there is an $r$ small enough to ensure that $B_{S}(c ; r)$ is a topological disk.

For a sufficiently dense sample set $P$, the iDt of $P \subset S$ can be defined analogously to the Euclidean Delaunay triangulation. We demand that each triangle have an empty geodesic circumdisk: a unique geodesic disk that is empty of all sample points. This yields the same combinatorial structure as we obtained via the nerve of the iVd.

Sampling criteria that guarantee that the iDt is well defined were proposed by Leibon and Letscher [LL00, Lei99]. A loose estimate on the actual number of samples this implies was published in [OI03]. In Chapter 5 we relax the criteria of Leibon and Letscher. It turns

[^3]

Figure 2.12: In a restricted Voronoi diagram the Voronoi cells need not be connected, as shown by the restriction to the red curve of the yellow Voronoi cell near the centre of the figure. The Voronoi cells of the intrinsic Voronoi diagram by contrast are always connected.
out that the density requirements that ensure a well formed iVd are also sufficient to ensure that there is a unique geodesic of minimal length between each pair of neighbouring samples in $P$.

The iDt and the iDt-mesh have not received the same attention in geometry processing as has the rDt. This is presumably because using the Euclidean distance $d_{\mathbb{R}^{3}}(p, q)$ to measure the distance between two points $p, q \in S$ seems simple and convenient in practice. However, there are advantages to using the metric of geodesic distance, $d_{S}(p, q)$.

We have the relation $d_{\mathbb{R}^{3}}(p, q) \leq d_{S}(p, q)$. The difference between $d_{\mathbb{R}^{3}}(p, q)$ and $d_{S}(p, q)$ is most pronounced when $p$ and $q$ lie close in the ambient space but far on the surface, as would be the case if we chose $p$ and $q$ to lie at the centre, but on opposite sides on the surface of a pancake, for example. (See Figure 2.12.)

The Voronoi cells of the intrinsic Voronoi diagram are connected in the topological sense (in fact they are path connected). This follows from the following observation, which Aurenhammer and Klein [AK00] make in the context of nice metrics:

Lemma 2.6 If $a \in \mathcal{V}(p)$ and $\gamma$ is a minimal geodesic between $a$ and $p$, then $\gamma$ is contained in $\mathcal{V}(p)$.

Proof Suppose to the contrary that $\exists w \in \gamma$ and $q \in P$ with $q \neq p$ and $w \in \mathcal{V}(q)$ but $w \notin \mathcal{V}(p)$. It follows that $d_{S}(a, q) \leq d_{S}(a, w)+d_{S}(w, q)<d_{S}(a, w)+d_{S}(w, p)=d_{S}(a, p)$, contradicting $a \in \mathcal{V}(p)$.

In contrast, the Voronoi cells of the rVd may contain more than one connected component. An example may be easily constructed using the pancake mentioned above. Thus the intrinsic metric is in a sense "nicer" than the Euclidean metric in the context of surfaces embedded in $\mathbb{R}^{3}$.

Another argument in favour of the intrinsic metric is that we are liberated from dependence on the ambient space. This argument is weak in the context of computer graphics where an attempt to divorce the surface from the surrounding space appears to be an exercise in senseless abstraction. However, for many other applications, especially when the surface lies in a higher dimensional space, the intrinsic metric is the only sensible choice. This is because working in the high dimensional ambient space renders one subject to "the curse of dimensionality": some computations become intractable in a high dimensional setting. Finally, an analysis based on the intrinsic metric allows us to tap into a wealth of results in Riemannian geometry, and in some cases stronger statements can be made in this context than have been made with more traditional extrinsic methods. The intrinsic sampling criteria we present in Chapter 5 is an example of this.

### 2.3.4 Related structures

The rDt and the iDt-mesh have both been well studied and proven themselves to be important structures in theory and practice. However, they share a trait that can be problematic in practice: for their definition, both the rDt and the iDt-mesh depend in an essential way on the underlying surface $S$ that is being approximated. This is an obvious difficulty in surface reconstruction, for example, where $S$ is known only by the samples $P$, and some assumptions on the density of $P$ and on the regularity of $S$. In this case, when one has constructed a mesh $M$ whose vertices are $P$, one would like to make claims of the fitness of $M$ as a representation of $S$. If it could be demonstrated that $M$ is an rDt or an iDtmesh, then an appeal could be made to the known qualities of these structures as surface
approximations (results which we will discuss in Section 2.4). Unfortunately, such a direct approach is impossible since $S$ is unknown.

In general, given a mesh $M$, there are no known algorithms which enable one to decide if $M$ is an rDt or an iDt-mesh of some smooth surface $S$. As we will see in Chapter 3, such an algorithm is unattainable in principle. Thus there is a motivation to define alternate mesh structures which display good approximation properties and which can be identified algorithmically based on the properties of $M$ alone. We present here several mesh structures which exploit the Delaunay paradigm to some degree. Many of these structures do not depend upon a reference surface $S$ for their definition, but nonetheless have proven approximation guarantees. These structures have arisen mostly within the context of surface reconstruction.

We focus on the structures themselves, not on the algorithms which seek to produce them. For comprehensive surveys of Delaunay-based surface reconstruction algorithms we refer to the article by Cazals and Giesen [CG06], and the book by Dey [Dey07].

## Cocone meshes

Many surface reconstruction algorithms provide guarantees on the quality of the output surface if specific sampling density assumptions are met. The first such algorithm was the Crust algorithm developed by Amenta and Bern [AB98]. This was a seminal paper in which several important concepts were introduced including the local feature size and the pole of a sample point. The pole of a sample $p$ is the Voronoi vertex of $V(p)$ that is farthest from $p$ (for unbounded Voronoi cells the pole is taken as a point at infinity in the direction of the average of the unbounded Voronoi edges). The pole vector is the normalized vector, $\hat{\boldsymbol{v}}_{p}$, from $p$ to its pole, and it provides a good estimate of $\hat{\boldsymbol{n}}_{S}(p)$. The mesh $M$ is extracted from the Delaunay tetrahedralization of the poles and the original samples together; only triangles whose vertices are all samples are considered while constructing $M$.

The algorithm and the analysis was subsequently simplified with the Co-cone algorithm [ACDL00]. Again the pole vectors are used to estimate the surface normals at the sample points. At each sample point a cocone is conceptually constructed. The cocone at $p$ is a portion of $V(p)$ that is the complement to a fat double cone whose axis is the pole vector at $p$. Specifically, the cocone at $p$ is given by $\left\{x \in V(p) \left\lvert\, \angle_{a}\left(\overrightarrow{p x}, \hat{\boldsymbol{v}}_{p}\right) \geq \frac{3 \pi}{8}\right.\right.$, where $\angle_{a}(\boldsymbol{a}, \boldsymbol{b})$ is the acute angle between the linear spaces generated by $\boldsymbol{a}$ and $\boldsymbol{b}$.

Candidate triangles for the reconstruction are selected from $\mathcal{D}^{3}(P)$. A Delaunay triangle
$t$ is a cocone triangle if its dual Voronoi edge meets the cocone associated with each of its three vertices. The final stage in the reconstruction selects a manifold mesh from the cocone triangles. This is always possible, if the specified sampling criteria are met, because the cocone triangles contain all the triangles of the rDt. The resulting mesh is a cocone mesh; a closed manifold substructure of $\mathcal{D}^{3}(P)$, all of whose faces are cocone triangles. It is shown that the cocone meshes are good approximations to $S$.

The 3D Delaunay tetrahedralization and the rDt played prominently in the design and analysis of these algorithms, and it is fair to say that the scaffold provided by the Delaunay tetrahedralization has played a role in most subsequent reconstruction and meshing algorithms that have provable guarantees on their output quality.

## Regular interpolants

Although the definition of a cocone mesh does not invoke the unknown surface $S$, the correctness of a cocone mesh as a representative of $S$ can only assured if the required sampling criteria are met. But these sampling criteria are defined explicitly with respect to properties of the unknown surface, as we will discuss in Section 2.4. This motivates a search for alternate sampling criteria which can be algorithmically verified based upon the properties of the reconstruction alone.

In a notable work in this direction Petitjean and Boyer [PB01] introduced the notion of regular interpolants. They defined these by means of the discrete medial axis of a mesh $M$. This is a subset of the 3D Voronoi diagram of the vertices of $M$ : It is the union of all of the Voronoi edges, facets and vertices which are not dual to a simplex of $M$. They then defined the local thickness of $M$ at a vertex $p$, which is essentially a measure of the distance from $p$ to the discrete medial axis. Then $M$ is a regular interpolant if for each $p \in M$ the local thickness is larger than the circumradius of each triangle in $M$ that is incident to $p$. Although the definition is attractive in its simplicity, it suffers a fatal flaw. The problem is the existence of sliver tetrahedra in $\mathcal{D}^{3}(P)$.

It was pointed out by Amenta and Bern [AB98] that sliver tetrahedra have dual Voronoi vertices that may be arbitrarily close to the original surface $S$. Sliver tetrahedra are those whose vertices are distributed near the equator of their circumsphere. They are flat and have a small volume, but no short edges (See Figure 2.13) ${ }^{6}$. Such tetrahedra cannot be

[^4]

Figure 2.13: A sliver tetrahedron in $\mathcal{D}^{3}(P)$ can have its circumcentre (Voronoi vertex) arbitrarily close to the surface. The discrete medial axis defined in [PB01] includes only those Voronoi edges and facets that are not dual to faces or edges of the mesh, but it includes all the Voronoi vertices. In (a), a sliver tetrahedron is depicted in its circumsphere, the centre of which is a Voronoi vertex (blue) near the surface. In (b) the Voronoi cell of the far sample is depicted, with the green curve representing where it intersects the surface. The blue Voronoi vertex is connected to three edges of this cell, but not all of them can be dual to faces of a single smooth manifold mesh.
eliminated by simply increasing the sampling density and as a consequence, it follows that the discrete medial axis may also be arbitrarily close to $S$, regardless of the sampling density. Thus regular interpolants are too restrictive; we are not guaranteed to achieve a regular interpolant for any given density based sampling criteria.

In a simultaneously published paper, Dey and Giesen [DG01] demonstrated that the local feature size sampling constraint required by the cocone algorithm can indeed be verified based on the properties of the Voronoi diagram of the point set alone, and regular interpolants have since drawn little interest.

## Alpha, Beta, and Gabriel complexes

Although regular interpolants themselves have not drawn much attention, the work did introduce a structure which has reappeared in several subsequent publications. Petitjean
and Boyer [PB01] showed that a regular interpolant must be a substructure of the Gabriel complex, which they also introduced in that paper. They defined the Gabriel complex as those triangles in $D t$ whose circumscribing ball defined by the circumcentre and circumradius of the triangle, is empty of all points of $P$. It is a natural extension of the Gabriel graph. We refer to a mesh that is a substructure of this complex as a Gabriel mesh, and we will study these structures in detail in Chapter 4.

A related complex is the $\alpha$-shape [EM94]. This consists of those triangles that admit an empty ball of radius $\alpha$ on their dual Voronoi edge. Surface reconstruction algorithms which are based on the $\alpha$-shape generally need to assume a uniform sampling density.

The $\beta$-skeleton [KR85] was exploited for curve reconstruction [ABE98], but it has not been generalized for the purposes of surface reconstruction. The $\beta$-skeleton consists of those edges in the planar Delaunay triangulation which possess empty disks that intersect at the endpoints of the edge and with radius of $\beta$ times the length of the edge.

## Polyhedra of minimal surface area

Among the early publications dealing with the surface reconstruction problem was a short paper by O'Rourke [O'R81] in which local edge swapping is performed to reduce the surface area of an initial mesh. The assumption was that the mesh of minimal surface area connecting the vertices $P$ would be one of the best representations of $S$ amongst all the possible ways of connecting $P$. It was mentioned that attaining the true minimal surface area mesh would probably not be computationally practical, since the lower dimensional analogous problem of finding a minimal length polyhedron connecting a set of points in the plane is essentially the travelling salesman problem, which is known to be NP complete.

However, the travelling salesman problem is NP-complete for a generic set of points, but points which are densely sampled from some smooth curve have properties which reflect their special origin. In a series of papers initiated by Giesen [Gie99a, Gie99b, AM01], it was demonstrated that the travelling salesman perspective yields a powerful algorithm for producing a polygonal reconstruction of the original planar curve. These algorithms are able to perform reliably on much weaker regularity assumptions than are required by Voronoibased curve reconstruction algorithms.

For the case of surface reconstruction, an algorithm that seeks to produce a polyhedron of minimal surface area was introduced by Althaus and Fink [AF02]. However the algorithm applies only to data sampled on parallel planar contours. A contour on $S$ is defined
as the curve of intersection between $S$ and a plane in $\mathbb{R}^{3}$. The contours parallel if their corresponding planes are parallel.

Although Alboul et al. [AKTvD00] suggest that minimizing surface area does not produce a desirable mesh in general, the example they give involves a point set that manifestly does not meet the sampling criteria that is demanded in surface reconstruction or isotropic surface remeshing algorithms. While a mesh with minimal surface area may not be the best representation of $S$ by most measures, there is no evidence to suggest that such a mesh will not be amongst the good ones, when $P$ is well sampled from $S$.

We mention these works because the connection with the Delaunay paradigm will become apparent as we pursue our investigation into Delaunay meshes. In particular, we will observe that an extrinsic Delaunay edge flip is one which reduces the surface area of the mesh. We will see that if it can be shown that a polyhedron of minimal surface area on $P$ is a smooth mesh, then it must be a Delaunay mesh.

## Flow complex and relatives

There have been works which present algorithms which are metaphorically referred to as shrink wrapping algorithms, and these give the impression of being related to surface area minimization. Among these were later works by Giesen et al. [GJ02, GJ03] which introduced the flow complex. The flow complex is a piecewise linear structure whose vertices are the critical points of the distance function as measured from the sample point set, as well as the sample points themselves. The precise description of the flow structure is quite involved. It is neither a substructure of the 3D Delaunay tetrahedralization nor of the Voronoi diagram, but it can be described completely in terms of them both. The wrap algorithm that was developed independently by Edelsbrunner [Ede04] is closely related to the flow complex, but it employs a power distance from Voronoi vertices. It can be seen as the dual of the flow complex, but the critical points are the same [RS07]. In contrast to the mesh constructed by the flow complex, the output of the wrap algorithm is a subcomplex of the 3D Delaunay tetrahedralization. Another notable shrink wrapping algorithm with this property was presented by Chaine [Cha03]. The algorithm is described as a convection algorithm and is perhaps the most worthy of the shrink wrapping classification. Interestingly this work draws connections with Gabriel meshes.

Although the shrink wrapping algorithms give the impression of surface area minimization, just as their metaphorical moniker implies, no explicit claims of surface area minimization are made. These algorithms have not been analyzed from this perspective.

## Alternate metrics

The same Delaunay approach that is employed to produce the rDt and the iDt-mesh, can be used to define other mesh structures by simply defining an alternate metric on $S$.

In adaptive sampling theory, the density of the sample set $P$ on $S$ is governed by some sizing function, which typically specifies a higher sampling density in regions of higher curvature. Various sizing functions are mentioned in Section 2.4, and the relationship between them is examined in Chapter 5. In some contexts it is natural to want to modulate the intrinsic distance function on $S$ by an appropriate sizing function. A nice example of where exactly such an approach is used to produce a Delaunay structure is in the farthest point sampling algorithm presented by Peyré and Cohen [PC03, PC06]. In this algorithm the Voronoi diagram of $P$ with respect to the modified metric is maintained while iteratively adding samples at points on $S$ that are farthest from the current $P$ with respect to this metric. The final mesh that is produced is the dual to the resulting Voronoi diagram.

A common meshing technique is to define the connectivity locally via the Delaunay triangulation on a local planar parameterization, the works of Chen and Bishop [CB97], and Gopi et al. [GKS00] being two notable examples. This method can be loosely described as an alternate metric technique in that the geodesic distances are being distorted by the parameterization. Chen and Bishop actually employ Delaunay refinement using an alternate metric in the plane so as to minimize the distortion induced by the parameterization.

Approaching the modified metric from a more discrete perspective, Glickenstein [Gli05] studied weighted triangulations of piecewise flat surfaces. The corresponding Delaunay structure is a regular triangulation, which is the dual of the power diagram and these are defined just as in the Euclidean setting. Such structures have been studied extensively in Euclidean domains [AK00], [Ede01], but the systematic study of them in the context of surface meshes is in its infancy.

## Well centred meshes

Finally, there has been recent interest in imposing a stronger criterion than the Delaunay condition on a triangle mesh. Specifically, one may demand that no triangle in the mesh contains an obtuse angle. Such a nonobtuse mesh is necessarily a self-Delaunay mesh, however in general a self-Delaunay mesh will contain obtuse angles. For a given sample set $P$, a non obtuse mesh interpolating $P$ will not exist in general, however approximation algorithms exist which will produce a nonobtuse mesh with small perturbations of the original samples [LZ06]. Such meshes are convenient for certain algorithms, such as the fast marching algorithm of Kimmel and Sethian [KS98], which become much more complicated in the presence of obtuse angles.

The idea of a nonobtuse mesh extends to higher dimensions where the resulting simplicial complexes are referred to as well centred meshes. This is an allusion to the fact that the circumcentre of each simplex must reside within the simplex itself. Algorithms for producing such structures have recently been proposed [VHGR08]. In this context the interest in well centred meshes stems from the theory of discrete exterior calculus, where the theory is simplified if such a mesh can be assumed. In particular, the orthogonal dual structure to the mesh [Gli05] will be free of facets with negative volume. However, Glickenstein [Gli05] has shown that a weaker condition suffices to supply this simplification; it is sufficient to ensure that the dual facets of dimension one have positive length. For two dimensional domains with unweighted vertices, it is enough that the mesh be a self-Delaunay mesh.

### 2.4 Surface approximation theory

Having reviewed the principle triangle mesh structures that have been employed to represent surfaces, we now turn our attention to the question of what constitutes a good representation of a surface. Of course the answer to such a question depends very much upon the intended application of the structure.

The principle concern in many, if not most applications is the accuracy of the geometric approximation. The quality of geometric approximation can be evaluated based upon a three level hierarchy:

1. topological consistency
2. pointwise approximation
3. normal approximation

At the first level, we require for topological consistency that the approximation $M$ be homeomorphic to the original surface $S$. In other words we demand that $M$ and $S$ both have the same genus. In this sense the first level is of a different nature from level (2) and (3) in that it is either attained or it is not. The latter geometric criteria can be viewed as convergence issues where the quality of the approximation is expected to improve with increased sampling density.

Normally, in order to evaluate the quality of the pointwise and normal approximation of $M$, one assumes that the criterion of topological consistency has been met. This then allows for the definition of a homeomorphism, a bijective, Bi-continuous mapping: $\xi: M \rightarrow$ $S$. Then the pointwise approximation error of $M$ can be evaluated as some measure of the difference between $x$ and $\xi(x)$ for each $x \in M$. Likewise, the quality of the normal approximation is judged by a comparison of the normal to $M$ at $x$ with the normal to $S$ at $\xi(x)$.

A family of meshes will be useful as representations of the geometry of $S$ if it can be shown that they will converge to $S$ in both position and normal as the number of vertices is increased. The two types of convergence need to be evaluated independently. A classic example of a family of meshes that converges pointwise to a cylinder, but which display no normal convergence (hence no convergence in surface area), is the Schwarz lantern [HPW06] (see Figure 2.14).

The meshing problem involves not only sampling $S$, but also deciding how those samples are connected together to form the triangle mesh. The Delaunay paradigm has been exploited in both these aspects of meshing. The connectivity of the samples will influence the geometric accuracy of the mesh approximation. However, as we discuss in Section 2.4.3, another criterion that is influenced by connectivity is element quality: how close are the triangle faces to equilateral triangles. Optimizing for geometric accuracy alone may lead to meshes with poor element quality. This is undesirable in some applications.

### 2.4.1 Sampling and topological consistency

Perhaps the most fundamental issue that must be addressed when meshing or reconstructing a surface is determining the distribution of samples necessary to ensure that the final mesh is topologically equivalent to the original surface. Ideally the sampling density is specified by


Figure 2.14: The Schwarz lantern is a triangle mesh of a cylinder that does not display normal convergence as the sampling increases. (a) Here we have used the same tessellation pattern on a torus. (b) A self-Delaunay mesh on the same vertex set. The triangle normals are more consistent with the expected surface normals.
an adaptive sampling criterion. Such criteria generally impose restrictions on the sampling density based on local curvature properties as well as semi-local properties relating to some notion of the distance to "the other side" of the surface. We refer to functions that can be used to modulate the sampling density in this way as sizing functions. These functions take positive values which can be thought of as having the units of distance. Thus at each point on the surface, a sizing function can be used to specify a radius within which a certain proportion of representative samples is expected. The square of such a function can, for example, be used to define a weighted area measure for governing stochastic sampling.

The best known sizing function is the local feature size (lfs), introduced by Amenta and Bern [AB98]. Its description requires some preliminary definitions ${ }^{7}$.

Given a closed set $C \subset \mathbb{R}^{3}$, e.g., a surface, a medial ball is an open ball $B \in \mathbb{R}^{3} \backslash C$ that is maximal with respect to inclusion (i.e., no other open ball in $\mathbb{R}^{3} \backslash C$ contains $B$ ). If $p \in C$ is contained in the closure of $B$, we say $B$ is a medial ball at $p$. The medial axis of $C$ is the

[^5]

Figure 2.15: The medial axis is given by the closure of the set of centres of medial balls. The medial balls are maximal with respect to inclusion in the complement of $S$. In other words a medial ball is defined by the property that it cannot be contained in a larger ball that does not intersect $S$.
closure of the set of centres of all the medial balls (see Figure 2.15). In the case of interest to us, $C$ is a smooth surface, $S$. In this case each point $x \in S$ will be associated with two medial balls, one on each side of the surface (one of them may have infinite radius), and these balls will be tangent to $S$ at $x$. The medial axis can be equivalently defined as the closure of the set of points $m \in \mathbb{R}^{3} \backslash S$ whose distance to $S$ is realized by more than one point in $S$. In other words, most medial balls touch $S$ in more than one point. However, at points where a principle curvature attains a local extrema, a medial ball may be an osculating ball which touches $S$ at only that point.

At a point $x \in S$, the lfs at $x$, denoted $\rho_{f}(x)$, is given by the Euclidean distance from $x$ to the medial axis of $S$. The local feature size becomes smaller in the presence of higher curvature. This is a property that is expected in any sizing function. Another property possessed by the lfs is that it is Lipschitz continuous, specifically it is a 1-Lipschitz function, which means that for any $p, q \in S,\left|\rho_{f}(p)-\rho_{f}(q)\right| \leq d_{\mathbb{R}^{3}}(p, q)$. This is especially convenient
for analysis because it enables us to bound $\rho_{f}(x)$ in a neighbourhood of a point $p$ if we have a bound on $\rho_{f}(p)$.

A density-based sampling criterion is given by the specification of a sampling radius, a function $\rho(x): S \rightarrow \mathbb{R}^{+}$such that for any point $x \in S$ there must be a $p \in P$ with $d(x, p)<\rho(x)$, where $\mathbb{R}^{+}$is the positive real numbers and $d$ is some metric on $S$.

In terms of the lfs, a typical sampling criterion demands that for any $x \in S$ there be a $p \in P$ such that $p \in B_{\mathbb{R}^{3}}\left(x ; \epsilon \rho_{f}(x)\right)$, the Euclidean ball of radius $\epsilon \rho_{f}(x)$ centred at $x$. That is, no point $x$ on the surface is farther than $\epsilon \rho_{f}(x)$ from a sample point, where $0<\epsilon \leq 1$. Thus $\epsilon \rho_{f}(x)$ is the sampling radius specified by the sampling criterion. A sample set that satisfies this criterion is called a lfs $\epsilon$-sample set for $S$. Amenta and Bern [AB98] demonstrated that the rDt will be homeomorphic to $S$ provided that $P$ fulfills such a sampling criteria, with $\epsilon<0.1$. Their proof relied on the closed ball property result of Edelsbrunner and Shah [ES94].


The majority of subsequent surface reconstruction and meshing algorithms with correctness guarantees have relied on the local feature size to define their sampling criteria. However, even when meshing, where $S$ is known, computing the lfs is problematic in practice since it requires knowledge of the medial axis, and estimates of the medial axis require a discretization of $S$, which is what is being sought in the first place. Cheng et al. [CDRR04] present a meshing algorithm with guaranteed topological consistency which avoids explicit reference to the local feature size. The algorithm is driven by violations of the closed ball property, and it strives for a sparse sampling, but there is no concise expression for the final sampling density.

There are other notions of topological consistency, besides demanding that $M$ be homeomorphic to $S$. One may require that they are ambient isotopic. This is a stronger requirement than homeomorphism. It essentially requires that the ambient space can be continuously deformed to bring $M$ to coincide with $S$, but without inducing any self-intersections along the way. As an example, a tube tied in a trefoil knot (Figure 2.16) is homeomorphic to a torus, but the two are not ambient isotopic. Amenta et al. [APR03] showed that if $M$ and $S$ are sufficiently close that a homeomorphism $\xi_{S}: M \rightarrow S$ is given by the projection to the closest point on $S$, and that for all $x \in M, d_{\mathbb{R}^{3}}\left(x, \xi_{S}(x)\right)<\frac{1}{2} \rho_{f}\left(\xi_{S}(x)\right)$,
then $M$ and $S$ are ambient isotopic. In light of that result, it is seen that most sampling criteria which guarantee that $M$ is homeomorphic to $S$, e.g. [AB98, Dey07], are sufficient to guarantee an ambient isotopy as well.

Another notion of topological consistency is the concept of homotopy equivalence. This concept is of little interest in the context of manifold surface approximation, but it becomes much more relevant when constructing a mesh to represent a non-manifold object. In this context, homotopy equivalence is a weaker criterion than a homeomorphism. Recent work by Gao et al. [GGOW08] provides guarantees of homotopy equivalence in the meshing of arbitrary planar domains.

The three point hierarchy presented in the introduction to this section is subject to debate. Specifically, it can be argued that topological consistency is not a prerequisite to a good geometric approximation. For closed surfaces, the topology is completely characterized by the genus: the number of handles (or holes) in the surface. From a topological perspective, all holes have equal importance, but from the point of view of geometry this is simply not true. For many applications it may not be important to preserve tiny handles. The notion of topological persistence was developed [ELZ00] as a way of capturing the relative size and importance of handles. Of course if topological equivalence is sacrificed, no homeomorphism between the original surface and its approximation is possible, so subsequent geometric accuracy analysis becomes more problematic.

### 2.4.2 Convergence and geometric accuracy

Assuming the criterion of topological consistency has been met, we turn our attention to the geometric fidelity of $M$ as a representation of $S$. In other words we are interested in evaluating the pointwise positional accuracy of $M$ and the accuracy of the face normals of the triangles of $M$ as representatives of corresponding surface normals on $S$.

## Pointwise convergence and the projection mapping

Even if the set $P$ of vertices of $M$ all lie on $S$ itself, there may be points of $M$ that are far from $S$ if the density of $P$ is not high enough with respect to the local feature size. If $d_{\mathbb{R}^{3}}(x, S)$ is the (minimum) distance from $x \in M$ to $S$, then one natural measure of the pointwise accuracy of $M$ is $\sup _{x \in M} d_{\mathbb{R}^{3}}(x, S)$. This is called the Hausdorff distance from $M$ to $S$. Notice that this distance is not symmetric: The Hausdorff distance from $S$ to $M$
will be different in general. The symmetric Hausdorff distance between $M$ and $S$ is the maximum of these two numbers. If we have a family of meshes $M_{n}$ whose vertices, $P$ are samples of $S$ with $|P|=n$, then if the symmetric Hausdorff distance between $M_{n}$ and $S$ goes to zero as $n$ goes to infinity, we say the family $\left\{M_{n}\right\}$ displays pointwise convergence.

If $M$ is pointwise sufficiently close to $S$, then there is a natural mapping $M \rightarrow S$ that embodies the concept of "distance to the closest point" implied in the above discussion. If $m \in \mathbb{R}^{3}$ is not a point on the medial axis of $S$, then there is a unique point $p \in S$ such that $d_{\mathbb{R}^{3}}(m, p)=d_{\mathbb{R}^{3}}(m, S)$. The point $m$ must lie on the line generated by the normal to $S$ at $p$, so $p$ can be viewed as the orthogonal projection of $m$ onto $S$. For a set $U \subset \mathbb{R}^{3}$, the projection mapping $\xi_{S}: U \rightarrow S$, takes each $m$ in $U$ to its closest point in $S$. If $\left.\xi_{S}\right|_{M}$ is injective, then it defines a homeomorphism of $M$ onto $S$. As mentioned above, the existence of this homeomorphism implies that $M$ and $S$ are ambient isotopic. It is standard in geometric accuracy analysis to exploit this homeomorphism to demonstrate topological correctness.

For any smooth surface $S$, we can construct a neighbourhood $U \supset S$ upon which the projection $\xi_{S}$ is well defined (see also [dC76]). For a point $p \in S$, the local reach at $p$ is the radius of the smallest of the two medial balls at $p$. The local reach is a continuous function $\rho_{R}: S \rightarrow \mathbb{R}^{+}$[Fed59]. For each $p \in S$, let $I_{p}$ denote the open interval on the normal line through $p$, centred at $p$ and with length $2 \rho_{R}(p)$. By construction, any point $m \in I_{p}$ is closer to $p$ than to any other point in $S$, thus for $p \neq q, I_{p} \cap I_{q}=\emptyset$. The set $U_{\rho_{R}}=\cup_{p \in S} I_{p}$ is called a tubular neighbourhood of $S$, and the projection $\xi_{S}$ onto $S$ is well defined on $U_{\rho_{R}}$. In fact, $\xi_{S}$ is well defined on any neighbourhood of $S$ that does not contain any point of the medial axis.

## Normal convergence

Establishing pointwise convergence is not sufficient to guarantee that $M_{n}$ will be a good representation of $S$ for sufficiently high $n$. It is crucial that one also demonstrates normal convergence. To evaluate how effectively $M$ approximates the normals of $S$ one needs a means of corresponding points on $M$ with points on $S$. Typically the projection mapping is used, so are interested in the size of $\left\|\hat{\boldsymbol{n}}_{t}-\hat{\boldsymbol{n}}_{S}\left(\xi_{S}(x)\right)\right\|$, where $t$ is a triangle in $M, \hat{\boldsymbol{n}}_{t}$ is the normal to $t$, and $x$ is a point in $t$. If an orientation can be chosen on each $M_{n}$ so that the supremum of this number over all $t \in M_{n}$ and $x \in t$ goes to zero as $n$ goes to infinity, we say that $\left\{M_{n}\right\}$ displays normal convergence to $S$.

If a sample point $p$ is a vertex in $M_{n}$ for all $n$, then the normals to triangles that have $p$ as a vertex must converge to $\hat{\boldsymbol{n}}_{S}(p)$. This implies that the difference in the normals of two adjacent triangles must go to zero as $n$ goes to infinity. Thus any family of meshes which displays normal convergence will consist of smooth meshes, at least for sufficiently large $n$.

The normal vector at a vertex in a triangle mesh does not need to be defined. However, in computer graphics, defining a normal vector at the vertex of a mesh can be important. Most schemes for defining a normal at a vertex involve a weighted sum of the face normals of the adjacent triangles. Recent work indicates that weighting the face normals by the face angle at the vertex yields convenient properties [BA05]. In the emerging formalism of the discrete exterior calculus, the formulation of a consistent scheme for defining tangent planes at vertices is considered an open problem [DHLM05]. Any sane definition of a normal vector at the vertices will converge to the normal of the surface at that point if the meshes display normal convergence as we have defined it.

It has been shown [MT04], [HPW06] that normal convergence is equivalent to convergence of surface area and to the convergence of the intrinsic distance functions of the $M_{n}$ to that of $S$.

## Convergence properties of Delaunay structures

One of the attractive properties of both the rDt and the iDt-mesh is that these families both display pointwise and normal convergence. This was demonstrated for the former in [AB98], and more recently for the latter in [DLYG06, DLJ $\left.{ }^{+} 07\right]$. The bounds for the normal error that appeared in [AB98] as well as in many subsequent works, e.g., [ACDL00, Dey07], rested on a "normal variation lemma" which bounds the difference in the normals between two points whose separation is bounded by the lfs. A recent erratum, [AD07], has corrected the proof of that lemma and improved its bound. Thus the normal error bounds mentioned in affected previous works can be improved.

For general meshes, increasing sampling density alone is not sufficient to ensure a decrease in the normal error, as the Schwarz lantern will attest. However, if the circumradius of the triangles can be controlled, then so too can be the normal error. For general meshes, the normal error depends linearly on the largest circumradius of the triangles [MT04]. This highlights one of the reasons for the success of the Delaunay paradigm in surface meshing: Delaunay triangulations favour triangles with small circumradius. Actually, that statement merits further exploration. While it is known that the 2D planar Delaunay triangulation
minimizes the maximum circumradius of the triangles [DS89], we do not make corresponding claims for the structures presented in Section 2.3. It is unknown whether a mesh that minimizes the maximum circumradius, of all meshes homeomorphic to $S$ and with vertex set $P$, need be Delaunay in any sense.

However, the fact that Delaunay structures do constrain the triangle circumradius with respect to the sampling density can be seen if we take the approach of considering the dual Voronoi diagram on $S$ that results from the appropriate choice of metric $d$. If a sampling density is specified in terms of a sufficiently small sampling radius, $\rho$, then the point $z$ on $S$ that is farthest from any sample point will be a Voronoi vertex. Since the distance from $z$ to the closest sample points must be less than $\rho(z)$, the dual triangle to $z$ in the Delaunay structure will necessarily have a circumradius smaller than $\rho(z)$. Thus the Delaunay structure will display normal convergence as the sampling radius is reduced.

It is worth being more explicit on this point. Convenient bounds on the normal error and the positional error for a Euclidean triangle $t$ whose vertices belong to $S$ are given in [DLYG06][Theorem 2]. Specifically, if the vertices of $t$ are contained in $B_{S}(c ; r)$, with $r<1 /(4 \kappa)$, where $\kappa$ is a bound on the maximal curvature in $B_{S}(c ; r)$, then for any point $m \in t$, we have $\left\|\hat{\boldsymbol{n}}_{S}\left(\xi_{S}(m)\right)-\hat{\boldsymbol{n}}_{t}\right\| \leq 4.5 \kappa r$ and also $d_{\mathbb{R}^{3}}\left(m, \xi_{S}(m)\right) \leq 9 \kappa r^{2}$. Now if we have a sampling radius that demands $r<\epsilon \rho_{f}(c)$, for a sufficiently small $\epsilon$, then, using the Lipschitz continuity of the lfs and the relationship between maximal curvature and lfs discussed in Chapter 5, the normal error is bounded by $\mathcal{O}(\epsilon)$ and the positional error is bounded by $\mathcal{O}\left(\rho_{f}(c) \epsilon^{2}\right)$. It is possible to obtain the same bounds if we are using the Euclidean metric instead of the intrinsic metric (see Lemma 5.21 in Chapter 5). Thus these bounds hold for both the rDt and the iDt-mesh.

### 2.4.3 Element quality

There is another quality of $M$ that is often sought in addition to, or even in preference to geometric fidelity. We often want to ensure that $M$ contains no "bad triangles". This is the goal of good element quality. What exactly is meant by a bad triangle is the subject of a sizable paper by Shewchuk [She02]. There are many measures for judging the quality of a triangle, perhaps the best known is the ratio of the shortest edge to the circumradius, a number that is generally desired to be as large as possible. Good triangle quality is important for applications which discretize partial differential equations on $S$ [She02]. Computing eigenfunctions of the Laplace-Beltrami operator, a computation that is drawing increasing
interest in the geometry processing community [ZvKD10], can be seen to fall into this category.

Although there may be no clear consensus of exactly what defines a bad triangle, it is generally agreed that a triangle that is near to being equilateral is good. In general the Delaunay structures provide good element quality, but for a fixed sample budget, the best geometric representations of $S$ are composed of triangles that are far from equilateral. In the presence of anisotropic curvature, such as on the surface of a cylinder for example, long thin triangles oriented in the direction of low curvature provide a better geometric approximation than may be obtained with triangles that are closer to being equilateral.

A nice demonstration of this phenomenon is presented in works that seek to optimize triangle meshes by edge flipping algorithms geared towards minimizing some discrete curvature measure [DHKL01] [vDA95]. Although these works appeal primarily to the visual quality of the results and lack quantitative analysis, they are supported by more quantitative works: In [DLR90] the best planar mesh to represent a known, possibly anisotropic, function $f: D \subset \mathbb{R}^{2} \rightarrow \mathbb{R}$ is sought, and in [AKTvD00] experiments are performed to assess the geometric quality of the meshes produced by optimizing the quality measures discussed in [vDA95]. The meshes produced are visually appealing, but they contain many long skinny triangles. Interestingly, edge flipping algorithms that minimize the Willmore energy, which is, roughly speaking, the integral of the square of the mean curvature, tend to produce nice triangle quality at the expense of mesh smoothness [ABR06].

Thus it appears that the goal of geometric fidelity is at odds with the goal of good element quality. In particular, the Delaunay structures discussed in Section 2.3, despite the established convergence results for the rDt and the iDt-mesh, are not generally the best meshes for representing the geometry of $S$. This fact coupled with the demonstrated practical and theoretical attributes of Delaunay structures should motivate future work into Delaunay structures based on anisotropic metrics which reflect the principle curvatures.

Returning again to the question of what makes a bad element, we consider the issue from the point of view of geometric accuracy. There is evidence which supports the assertion that triangles with very large angles may be more problematic than triangles with very small angles. In the context of normal convergence, Morvan and Thibert [MT04] identified the rightness of a triangle $t$ as an important measure of element quality. The rightness of $t$ is simply the sine of the largest angle. The importance of rightness arises in the context of normal convergence of meshes approximating a smooth surface: the normal error is
minimized when the minimum rightness of the triangles is maximized. This result translates into the linear dependence of the normal error on the circumradius by the formula $r_{t}=\frac{\eta_{t}}{2 \gamma_{t}}$, where for triangle $t, r_{t}$ is its circumradius, $\eta_{t}$ is the length of its longest edge and $\gamma_{t}$ is its rightness.

There is an important hole in the approximation theory for surfaces. The data dependent triangulations produced by optimizing curvature energies tend to produce meshes with a smooth appearance, and there is some experimental evidence that these meshes have better approximation properties, in terms of positional error and normal error, than is obtained with the more isotropic triangulations produced by Delaunay methods. However, the resulting triangulations may have triangles with large circumradii, and even given the assumption of topological correctness, there are no theoretical geometric accuracy guarantees. In fact, all works which yield correctness guarantees on the approximation quality of a mesh family with respect to an unknown surface from which the vertices are sampled, rely on an assumption of triangle quality that essentially bounds the size of the circumradius. Without such a constraint, there is as yet no theoretical guarantee that a $\theta$-smooth mesh interpolating $P \subset S$ will be a good approximation of $S$, even for very small $\theta$.

## Chapter 3

## Self-Delaunay meshes

In this chapter we introduce self-Delaunay meshes, the structures central to this thesis. The natural place to begin the exposition is with the Delaunay triangulation of pwf surfaces. In the context of pwf surfaces, the theory of Delaunay triangulations enjoys a richness that is not found in intrinsic Delaunay triangulations of smooth surfaces. The Delaunay triangulation may not be a triangulation in the traditional sense. There may be more than one edge between two vertices, and a single edge may terminate at the same vertex at both ends. In Section 3.1 we examine this theory and demonstrate that the duality relationship with the intrinsic Voronoi diagram can be maintained.

In Section 3.2 we introduce self-Delaunay meshes themselves. We are interested in selfDelaunay meshes as models for smooth surfaces. This is also the primary role of the rDt and the iDt-mesh. If a smooth surface is densely sampled, then the circumradii of the triangles in all of these mesh structures will be small. If a self-Delaunay mesh is a good representation of $S$ and has sample points on $S$ as vertices, then the distance between sample points on $S$ as measured geodesically on $S$, will be very similar to the distance between these points as measured geodesically on the self-Delaunay mesh. So it is natural to ask whether or not an iDt-mesh must be a self-Delaunay mesh if the sample density is high enough. Similarly, for points that are close together, the Euclidean distance between them will be close to the geodesic distance between them, whether the geodesics are measured on $S$ or on a mesh that closely approximates $S$. However, we show in Section 3.3 that these three structures are distinct in general, regardless of the sampling density.

The distinction arises from patterns of four sample points that are almost cocircular in a given metric. We take the view that the different mesh structures arise from different choices
of metric on $S$. No matter how well a mesh $M$ approximates $S$, some metric distortion arises in the transition from the geodesic metric of $S$ to the geodesic metric of $M$. We can always have a sample point close enough to a circumcircle that it will be inside the circle when one metric is used, but outside in the other.

Such configurations may be viewed as being "almost degenerate", and the complaint could be made that the issue is one of splitting hairs. However, the problem is real, and it is closely related to the problem of sliver tetrahedra in the 3D Delaunay tetrahedralization. This is an unavoidable complication that must be faced by all Delaunay based surface reconstruction algorithms. Much of the research pursued in this thesis can be viewed as ultimately stemming from this issue.

We finish the chapter in Section 3.4 with a discussion of the advantages and drawbacks of the self-Delaunay mesh structure as compared with other Delaunay structures for surface representation.

### 3.1 Delaunay and Voronoi structures on pwf Surfaces

In 2005 Bobenko and Springborn presented a formal study and uniqueness proof of for the intrinsic Delaunay triangulation of the vertices of a piecewise flat surface [BS05] (formally published in [BS07]). Although the structure had been studied previously [Riv90, ILTC01], Bobenko and Springborn's work was the first to recognize its potential utility in geometry processing.

The definition of a triangulation employed in this context is weaker than the one we introduced in Section 2.2.3. The triangulation does not necessarily correspond to a simplicial complex. In particular, there may be more than one edge connecting two vertices, and an edge may be a loop that terminates at the same vertex at either end. Figure 3.2(a) on page 65 gives an example of such a triangulation of a tetrahedron, where the triangulation is depicted with red and black edges, and the green edges of the tetrahedron do not belong to the triangulation. These kinds of triangulations are natural in the context of pwf surfaces [Gli05]. In Hatcher's introduction to algebraic topology [Hat02], such triangulations are called $\Delta$-complexes. When the triangulation can be described by a simplicial complex, as in Section 2.2 .3 we say that it is a proper triangulation of the pwf surface.

The interest in Delaunay triangulations of pwf surfaces springs from Bobenko and

Springborn's observation [BS07] that the linear finite element discretization of the LaplaceBeltrami operator (the cotan operator [MDSB03]) has no negative edge weights on a Delaunay triangulation. Like all discrete Laplacian operators, the cotan operator can be described as a linear operator, L, which acts on functions defined at the vertices of the mesh such that for any such function $f$, and any vertex $p \in M, \mathrm{~L} f(p)$ is the difference between the value of $f$ at $p$ and a weighted sum of the values of $f$ at the one-ring neighbours of $p$. For a given edge $e=[p, q]$, the associated weight is given by half the sum of the cotangents of the two angles subtended by $e$. This operator admits an elegant geometric interpretation in terms of circumcentric dual cells, which are described in Appendix A.

A Laplacian operator without negative edge weights is desirable for a number of reasons outlined by Wardetzky et al. [WMKG07], including the preservation of a discrete maximum principle (harmonic functions have no maxima or minima in the interior of their domain). It has also been experimentally shown that the condition number of the operator can be significantly improved in most cases when the intrinsic Delaunay triangulation is used [FSBS06]. This improvement is beneficial to any application that involves the numerical evaluation of elliptic PDEs on triangle mesh surfaces. Examples include parameterization [DMA02] and reaction diffusion textures [Tur91]. In the context of parameterization, positive edge weights are also essential for the application of Tutte's theorem [Tut63, Flo98] which guarantees an injective planar embedding. However, in this context a proper triangulation is also required.

The nice properties that the Delaunay triangulation imparts on the cotan operator are revealed in a wider context in a discrete formulation of exterior calculus [DHLM05], in which the cotan operator arises naturally. In this context, the Delaunay triangulation can be seen to ensure not only positive weights in the cotan operator, but also that the natural inner product for discrete functions, defined by means of the discrete Hodge dual, is positive definite. As discussed in Section A.3, a Delaunay triangulation is not necessary to ensure this positive definite property, but to date a weaker family of triangulations which will ensure this has not been described.

In this section we show that the empty circumdisk property of Delaunay triangulations can be used to establish a Voronoi-Delaunay duality on pwf surfaces. In general, the dual of the Voronoi diagram will not be a proper triangulation (simplicial complex). However, we observe that a proper Delaunay triangulation is guaranteed if and only if the Voronoi diagram is well formed. This is of interest because a mesh requires a simplicial complex for its definition.

### 3.1.1 Delaunay Triangulations on pwf Surfaces

In this section $\mathcal{M}$ is a compact pwf surface without boundary. The discrete set $P \subset \mathcal{M}$ includes all the cone points of $\mathcal{M}$. We refer to the elements of $P$ as vertices, emphasizing that the model pwf surface we have in mind is a mesh. A tessellation of $\mathcal{M}$ with respect to $P$ is as follows.

A cell complex on $\mathcal{M}$ is a finite set of points $V \subset \mathcal{M}$ together with a finite set of open simple curve segments $E$ whose endpoints are in $V$, and a finite set of open topological disks $F$ such that the elements of $V, E$ and $F$ are together pairwise disjoint and their combined union is $\mathcal{M}$. The one skeleton of the cell complex, $\mathcal{G}$, is the union of the elements in $V$ and E.

A tessellation of $\mathcal{M}$ is a cell complex with further geometric constraints. The vertex set is $P$, and the elements of $E$ are geodesics, called the edges of the tessellation. The elements $f_{i} \in F$ are called faces, and for each face there exists a continuous map $\phi_{i}: Q_{i} \rightarrow \bar{f}_{i}$, where $Q_{i}$ is a closed planar polygon and $\bar{f}_{i}$ is the closure of $f_{i}$. The map $\phi_{i}$ is an isometry on the interior of $Q_{i}$, and is such that vertices of $Q_{i}$ get mapped to elements of $P$ that lie on the boundary of $f_{i}$. If $Q_{i}$ is an $n$-gon, we call $\bar{f}_{i}$ an $n$-gon face, and in particular, if $Q_{i}$ is a triangle, then we also call $\bar{f}_{i}$ a triangle face. A triangulation is a tessellation in which all the faces are triangle faces.

Note that edges cannot cross in a tessellation. The mappings $\phi_{i}$ are not required to be injective on the boundary of $Q_{i}$. In particular, two edges of $Q_{i}$ may be mapped onto a single edge in $E$. Likewise, the restriction of $\phi_{i}$ to the vertices of $Q_{i}$ is not required to be injective.

The Delaunay tessellation of $\mathcal{M}$ is defined in terms of empty disks. An immersed empty disk is a continuous map $\phi: \bar{D} \rightarrow \mathcal{M}$, where $D$ is an open disk in $\mathbb{R}^{2}$ and $\bar{D}$ is the closure of $D$, such that the restriction $\left.\phi\right|_{D}$ is an isometric immersion (i.e., every $p \in D$ has a neighbourhood which is mapped isometrically) and $\phi(D) \cap P=\emptyset$ (i.e., $\phi(D)$ is empty of vertices). We can think of $\phi$ as wrapping $D$ on $\mathcal{M}$, but it may wrap around onto itself: $\phi$ is not injective in general. It should be emphasized that $\phi$ is defined on the closure of $D$ and that only the image of $D$ itself is required to be empty. Most of the time we are working with empty disks that have elements of $P$ on their boundary, so that $\phi^{-1}(P)$ is non-empty.

The image of an immersed empty disk is an empty geodesic disk, but immersed empty disks are more convenient to work with than geodesic disks since they allow us to work with ordinary disks in the plane, with the caveat that the mapping $\phi$ is not injective in
general. Since $\mathcal{M}$ is flat in a neighbourhood not containing cone points, we can always find an isometric immersion $\phi$ whose image is a given empty geodesic disk and if two immersions $\phi$ and $\phi^{\prime}$ have the same geodesic disk as their image, then there will be a planar isomorphism $T: \mathbb{R}^{2} \rightarrow \mathbb{R}^{2}$ such that $\phi=\phi^{\prime} \circ T$.

Thus working with immersed empty disks is really equivalent to working with geodesic disks. However, the former allows us to place $D$ wherever is convenient on the plane. In particular, we have the following useful lemma, whose proof is indicated in [BS07, Lemma $6]$.

Lemma 3.1 Suppose that $\phi: \bar{D} \rightarrow \mathcal{M}$ and $\phi^{\prime}: \overline{D^{\prime}} \rightarrow \mathcal{M}$ are two immersed empty disks with $\phi(D) \cap \phi^{\prime}\left(D^{\prime}\right) \neq \emptyset$. Then there exists a disk $\tilde{D}$ with $\tilde{D} \cap D \neq \emptyset$, an isometry $T: \mathbb{R}^{2} \rightarrow \mathbb{R}^{2}$ with $T(\tilde{D})=D^{\prime}$, and an isometric immersion $\hat{\phi}: \overline{D \cup \tilde{D}} \rightarrow \mathcal{M}$ such that $\left.\hat{\phi}\right|_{\bar{D}}=\phi$ and $\left.\hat{\phi}\right|_{\widetilde{D}}=\phi^{\prime} \circ T$.

The Delaunay tessellation of $P$ on $\mathcal{M}$ is defined by the immersed empty disks $\phi: \bar{D} \rightarrow \mathcal{M}$ such that $\phi^{-1}(P)$ is non-empty. If $\phi^{-1}(P)$ contains three or more points, then its convex hull, $\operatorname{conv}\left[\phi^{-1}(P)\right]$, is a polygon and its image under $\phi$ defines a face of the tessellation. If $\phi^{-1}(P)$ contains exactly two points, then the image of $\operatorname{conv}\left[\phi^{-1}(P)\right]$ under $\phi$ is an edge. It was established [BS07] that these faces and edges do indeed describe a tessellation, something that is not obvious a priori.

If a face contains more than three vertices, the diagonals of the face are not included in the tessellation. To obtain a Delaunay triangulation, we triangulate all non-triangular faces. A face of the Delaunay triangulation is still contained in an immersed empty disk, but there may be more than three vertices on the disk's boundary.

The vertices are in general position if there exists no empty disk with more than three vertices on its boundary. In this case the Delaunay tessellation is the unique Delaunay triangulation of the vertices. We refer to the Delaunay triangulation of $P$ on $\mathcal{M}$ as the intrinsic Delaunay triangulation (iDt). However, it is worth emphasizing the distinction with the iDt on a smooth surface. In the pwf case we demand that $P$ include all the cone points of $\mathcal{M}$, but there are no density requirements on $P$. The price paid for relaxing the density requirements is that the triangulation is not proper in general. In contrast, a smooth surface $S$, has no distinguished points, but the iDt is only defined when the density of $P$ is sufficient to ensure that a unique proper geodesic triangulation exists.

Now, consider an arbitrary triangulation $\mathcal{T}$ of the vertices of $\mathcal{M}$. Since the triangles are
empty of cone points, they are intrinsically planar. Given an edge $e$ of $\mathcal{T}$, we can map the two triangular faces adjacent to $e$ isometrically onto the plane forming a quadrilateral with $e$ as its diagonal. Just as in the planar case, edge $e$ is locally Delaunay if it is contained in a disk that does not have the other two vertices of the quadrilateral in its interior. This is conveniently characterized by the fact that the sum of the angles subtending $e$ must not exceed $\pi$. This is different from the immersed empty disk criteria which characterizes a (globally) Delaunay edge in that only two additional vertices of $\mathcal{M}$ are considered. As usual, $e$ is nlD if it is not locally Delaunay.

As in the planar case, the iDt can be obtained by systematically flipping the geodesic edges that are nlD [BS07, ILTC01]. Namely, an nlD edge $e$ is replaced by the edge $e^{\prime}$ that is the other diagonal (guaranteed to be locally Delaunay) of the quadrilateral defined by the triangles adjacent to $e$. This algorithm runs in $\mathcal{O}\left(n^{2}\right)$ time, $n$ being the number of vertices in the mesh. The proof described in [She97] holds without modification to the case of a fixed pwf surface. In fact, the similarities to the planar case carry over to most of the properties that are shown to be optimized by Lawson's planar flip algorithm. The proofs based on the local properties of a flip-quad are unchanged since these quadrilaterals are isometric to planar quadrilaterals. Thus, for example, of all triangulations on $\mathcal{M}$, the iDt maximizes the minimum angle, and minimizes the maximum circumradius, and minimizes the harmonic index.

### 3.1.2 Voronoi Diagrams on pwf Surfaces

In this section we examine the intrinsic Voronoi diagram of a pwf surface and its relationship with the Delaunay tessellation. Since $\mathcal{M}$ is compact, in analogy with the Hopf-Rinow theorem [dC76] for smooth surfaces, a minimal geodesic exists between any two points in $\mathcal{M}$ [AZ67].

The intrinsic Voronoi diagram of $P$ on $\mathcal{M}$ is defined exactly as in Definition 2.4. The Voronoi cell, of $p \in P$, is given by $\mathcal{V}(p)=\left\{x \in \mathcal{M} \mid d_{\mathcal{M}}(p, x) \leq d_{\mathcal{M}}(q, x), \forall q \in P\right\}$, where $d_{\mathcal{M}}$ is the intrinsic metric on $\mathcal{M}$.

A Voronoi vertex is a point $c \in \mathcal{M}$ that has three or more distinct geodesics realizing the minimum distance from $c$ to $P$. A Voronoi edge is a curve $\gamma$ terminating at Voronoi vertices and such that every point $x$ in the interior of $\gamma$ has exactly two geodesics realizing the minimum distance from $x$ to $P$. If both the minimal geodesics connect with a common vertex in $P, \gamma$ is called an internal Voronoi edge. The interior of $\mathcal{V}(p)$ consists of those
points in $\mathcal{V}(p)$ which do not lie on a Voronoi edge or vertex (i.e., we exclude the interior Voronoi edges from the topological interior of the Voronoi cell). A point $x$ in the interior of $\mathcal{V}(p)$ is characterized by having a unique minimal geodesic connecting it with $p$.

An equivalent view of Voronoi edges and vertices is via the immersed empty disk property: If $\phi: \bar{D} \rightarrow \mathcal{M}$ is an immersed empty disk with centre $c$ and with $\phi^{-1}(P)$ containing three or more points, then $\phi(c)$ is a Voronoi vertex. If $\phi^{-1}(P)$ contains exactly two points, $p$ and $q$, then $c$ lies on a Voronoi edge, and it is an internal edge if $\phi(p)=\phi(q)$.

According to this view each Voronoi vertex is associated with a face in the Delaunay tessellation via the immersed empty disk that defines them both. Thus there is a finite number of Voronoi vertices. However, a Voronoi vertex is not necessarily associated with distinct samples and a Voronoi edge may terminate at the same Voronoi vertex at both ends.


Figure 3.1: If two immersed empty disks define the same Delaunay edge (the image under $\hat{\phi}$ of $[p, q]$ ), then their centres both lie on a common Voronoi edge that contains the image of $\left[c, c^{\prime}\right]$ under $\hat{\phi}$ (Lemma 3.5). Furthermore, since $\hat{\phi}$ is locally an isometry, the Voronoi edge must be a geodesic.

Voronoi edges are geodesics between Voronoi vertices. To see this, let $\phi: \bar{D} \rightarrow \mathcal{M}$ be an immersed empty disk with $\{p, q\}=\phi^{-1}(P)$ and $c \in D$ the centre. Thus $\phi(c)$ lies on some Voronoi edge $\gamma$. Since there are only two vertices on the boundary of $\phi(D)$, we can find some $\epsilon$ and (exploiting Lemma 3.1) another immersed empty disk $\phi^{\prime}: \overline{D^{\prime}} \rightarrow \mathcal{M}$ with centre $c^{\prime}, d_{\mathbb{R}^{2}}\left(c, c^{\prime}\right)=\epsilon$, with $\{p, q\}=\phi^{\prime-1}(P)$ and such that $\phi^{\prime}(p)=\phi(p)$ and $\phi^{\prime}(q)=\phi(q)$. Further, we have a mapping $\hat{\phi}: \overline{D \cup D^{\prime}} \rightarrow \mathcal{M}$ such that $\left.\hat{\phi}\right|_{\bar{D}}=\phi$ and $\left.\hat{\phi}\right|_{\overline{D^{\prime}}}=\phi^{\prime}$. Now any
point $\tilde{c}$ on the line segment $\left[c, c^{\prime}\right]$ will be the centre of a disk $\tilde{D} \subset \overline{D \cup D^{\prime}}$ that has $p$ and $q$ on its boundary (see Figure 3.1). Thus we have an immersed empty disk $\left.\hat{\phi}\right|_{\tilde{D}}: \bar{D} \rightarrow \mathcal{M}$ that has $\phi(p)$ and $\phi(q)$ as the only points of $P$ on its boundary. In other words, $\hat{\phi}\left(\left[c, c^{\prime}\right]\right)$ lies on the Voronoi edge $\gamma$, and it is geodesic, since $\left[c, c^{\prime}\right]$ is a geodesic in the plane, and $\hat{\phi}$ is an isometric immersion.

Lemma 3.2 A Voronoi cell is topologically a disk if and only if it contains no internal edges.

Proof Let $x \in \mathcal{V}(p)$ and assume that there are two minimal length geodesics, $\alpha$ and $\beta$, connecting $p$ with $x$. Suppose that $\mathcal{V}(p)$ were a topological disk. Together $\alpha$ and $\beta$ define a closed curve contained in $\mathcal{V}(p)$. Let $U$ be the region bounded by $\alpha$ and $\beta$. Then there is an isometric embedding $\varphi: U \hookrightarrow \mathbb{R}^{2}$. But then $\varphi(U)$ would be a region in the plane bounded by two geodesics (line segments) between $\varphi(p)$ and $\varphi(x)$. Thus $U$ must be empty and $\alpha=\beta$.

Conversely, if $\mathcal{V}(p)$ is not a disk, let $\alpha$ be a loop in a nontrivial homotopy class in $\mathcal{V}(p)$ and with $\alpha(0)=\alpha(\ell(\alpha))=p$. Let

$$
s_{0}=\sup \left\{s \mid \forall t \leq s, \exists \gamma_{p \alpha(t)} \text { s.t. } \alpha([0, t]) \cup \gamma_{p \alpha(t)} \text { has trivial homotopy }\right\},
$$

where $\gamma_{p \alpha(t)}$ denotes a minimizing geodesic between $p$ and $\alpha(t)$. Then $\alpha\left(s_{0}\right)$ will lie on an internal Voronoi edge.

Indeed, consider the immersed empty disk $\phi: \bar{D} \rightarrow \mathcal{M}$ such that $\phi(c)=\alpha\left(s_{0}-\delta\right)$, where $c$ is the centre of $D, \delta$ is arbitrarily small, and $p \in \phi(\bar{D})$. The preimage of $\gamma_{p \alpha\left(s_{0}-\delta\right)}$ will include a line segment that is a radius of $D$ with endpoint $b \in \bar{D}$ such that $\phi(b)=\alpha(0)=p$. Also, we may assume that $\delta$ is small enough so that there is some disk $B=B_{\mathbb{R}^{2}}(c ; \epsilon)$ that is isometric to its image under $\phi$, and $\alpha\left(\left[s_{0}-\delta, s_{0}\right]\right) \subset \phi(B)$. Let $a \in B, \phi(a)=\alpha\left(s_{0}\right)$. Then the image of $[a, b]$ under $\phi$ must be a minimal geodesic, $\gamma_{p \alpha\left(s_{0}\right)}$, between $p$ and $\alpha\left(s_{0}\right)$. Also, since the image of the triangle $[a, b, c]$ is null-homotopic, $\gamma_{p \alpha\left(s_{0}\right)} \cup \alpha\left(\left[0, s_{0}\right]\right)$ also has trivial homotopy type.

In a similar manner, exploiting a point $\alpha\left(s_{0}+\delta\right)$, we construct a minimal geodesic between $p$ and $\alpha\left(s_{0}\right)$ that is not homotopic in $\mathcal{V}(p)$ to $\alpha\left(\left[0, s_{0}\right]\right)$. Thus $\alpha\left(s_{0}\right)$ admits two distinct minimizing geodesics to $p$, and therefore lies on an internal Voronoi edge.

Remark 3.3 Lemma 3.2 does not extend to smooth surfaces. Using the same definition of internal edges, a Voronoi cell on a smooth surface may be a topological disk and yet still
contain internal Voronoi edges. This will become clear with the development of conjugate points and the cut locus in Chapter 5. However, If $\mathcal{V}(p) \subset S$ is not a topological disk, then it must contain an internal Voronoi edge. Rather than relying on immersed empty disks, as above, we can adapt the proof of [Cha06][Theorem IV.5.1], to show that $\alpha\left(s_{0}\right)$ admits two distinct minimizing geodesics to $p$.

Since a minimal closed geodesic in $\mathcal{V}(p)$ must pass through an internal edge, the interior of $\mathcal{V}(p)$ - that part which remains when we remove all Voronoi edges - is a topological open disk. Note also that we cannot have a Voronoi edge that is a closed loop not containing any Voronoi vertices. If such a loop were to exist, it would have to be the unique boundary between two Voronoi cells that were both topologically disks (otherwise an internal or other edge would create a Voronoi vertex). Therefore $\mathcal{M}$ must have only two vertices and be topologically a sphere. If such a pwf surface exists, it certainly cannot be realized as a mesh and it will not concern us here. These observations demonstrate that the Voronoi diagram can be viewed as a cell complex, the faces of which are the interiors of the Voronoi cells.

We now turn our attention to the duality relationship between the Delaunay tessellation and the Voronoi diagram. A nice thing about pwf surfaces is that if $\phi: D \rightarrow \mathcal{M}$ is an immersed empty disk, and $\phi^{-1}(P)=\{p, q\}$, then there is a unique geodesic between $\phi(p)$ and $\phi(q)$ contained in the image of $\phi$; it is the image of the line segment between $p$ and $q$. In other words there is only one possible edge contained in an empty disk with two samples on its boundary.

Furthermore, the image of the centre of $D$ lies on a Voronoi edge $\gamma$. If $e=[\phi(p), \phi(q)]$ is the Delaunay edge defined by $\phi$, then we say $\gamma$ is the Voronoi edge associated with $e$ and vice versa. The following lemmas demonstrate that this association is exclusive.

Lemma 3.4 There is a unique Delaunay edge associated with each Voronoi edge.

Proof Suppose that $e=[a, b]$ and $e^{\prime}=[a, b]$ are two Delaunay edges associated with the Voronoi edge $\gamma$. Let $u$ and $u^{\prime}$ be the centres of empty geodesic disks containing $e$ and $e^{\prime}$, respectively, and defining the association with $\gamma$. Now centred at every point between $u$ and $u^{\prime}$ on $\gamma$ there is an empty immersed disk with $a$ and $b$ on its boundary. Two such disks, if their centres are sufficiently close to each other, must contain the same Delaunay edge. Thus we can push the disk centre from $u$ to $u^{\prime}$ while always keeping $e$ in the empty disk. As a result, we must have $e^{\prime}=e$.

Lemma 3.5 Different Voronoi edges are associated with distinct Delaunay edges.

Proof Let $e=[a, b]$ be a Delaunay edge and suppose that it is contained in two different empty immersed disks $\phi: \bar{D} \rightarrow \mathcal{M}$ and $\phi^{\prime}: \overline{D^{\prime}} \rightarrow \mathcal{M}$. By Lemma 3.1 we can assume that $D \cap D^{\prime}$ contains a line segment whose image under the combined map $\hat{\phi}$ is $e$. We have $p, q \in \partial D \cap \partial D^{\prime}$ with $\hat{\phi}(p)=a$ and $\hat{\phi}(q)=b$ (Figure 3.1). Let $c$ and $c^{\prime}$ be the centres of $D$ and $D^{\prime}$, respectively. Then centred at any point on the line between $c$ and $c^{\prime}$ there is a disk $\tilde{D}$ that is contained in $D \cup D^{\prime}$ and touching $p$ and $q$ on its boundary. The restriction of $\hat{\phi}$ to $\tilde{D}$ defines an immersed empty disk. Therefore there is no Voronoi vertex between $c$ and $c^{\prime}$ and thus they must lie on the same Voronoi edge.

These results are summarized in Theorem 3.6, establishing a Voronoi-Delaunay duality on pwf surfaces.

Theorem 3.6 Considered together with its internal edges, the Voronoi diagram of the vertices of a pwf surface is a tessellation. Further, the empty circumdisk property defines a one-to-one correspondence between the edges of the Voronoi diagram and the edges of the Delaunay tessellation.

Remark 3.7 Glickenstein [Gli05] developed an elegant framework of duality structures for triangulations on pwf surfaces (and higher dimensional analogues). The triangulations considered are as we defined them here; they are not proper in general. Every triangulation has an associated dual complex, and when the triangulation is Delaunay, the dual complex is the Voronoi diagram. Thus Glickenstein's work encompasses the observations made by Theorem 3.6, which we published in [DZM07b]. We became aware of this work subsequently. Our exposition does serve to highlight the role of internal Voronoi edges, which were never explicitly mentioned by Glickenstein.

### 3.1.3 Proper Triangulations

Ultimately, we seek a triangulation of $\mathcal{M}$ that will allow us to define a mesh. Therefore we need to equip $\mathcal{M}$ with a triangulation that will give it the structure of a simplicial complex: a proper triangulation.

When $\mathcal{M}$ is given to us as the underlying pwf surface of a mesh $M$, then it already comes equipped with a proper triangulation, defined by $\mathcal{K}_{M}$, which describes the mesh faces. We
call this triangulation the extrinsic triangulation in order to distinguish it from the iDt. Figure 3.2(a) gives an example of a Delaunay triangulation that is not proper, overlaid on the extrinsic triangulation.


Figure 3.2: An example of a non-proper intrinsic Delaunay triangulation of a pwf surface (a), and its dual Voronoi diagram (b). The surface here is a tetrahedron defined by the triangle faces $[a, b, c],[b, c, d],[a, c, d]$, and $[a, d, b]$. In (a), the intrinsic Delaunay triangulation is defined by the red and black geodesics. The black geodesics coincide with the original mesh edges, but the red geodesics traverse faces that define the tetrahedron. It is composed of triangle $[a, b, b]$, which has a loop edge and shares an edge with itself; triangle $[c, d, c]$, which has the same characteristics; and triangles $[b, c, b]$, and $[c, c, b]$, which share two edges and each have a loop edge. In (b), The Voronoi edges are blue and $\mathcal{V}(b)$ and $\mathcal{V}(c)$ have internal Voronoi edges that are dual to the loop Delaunay edges. Note also that $\mathcal{V}(b)$ and $\mathcal{V}(c)$ share two distinct Voronoi edges due to the Voronoi vertices created by the internal Voronoi edges. Thus there are two Delaunay edges between $b$ and $c$ in the iDt.

Having established the duality between the iVd and the iDt, we can describe the conditions that will ensure that the iDt is proper. Edelsbrunner and Shah's result [ES94] ensures that if the iVd is well formed, then $\mathcal{M}$ will be homeomorphic to a geometric realization of the nerve of the iVd . It follows that there exists a proper triangulation of $\mathcal{M}$ with the connectivity of the iDt. It is also illuminating to establish this result directly.

Theorem 3.8 Let $\mathcal{M}$ be a pwf surface with vertex set $P$. The intrinsic Delaunay triangulation of $P$ on $\mathcal{M}$ is proper if and only if the intrinsic Voronoi diagram of $P$ on $\mathcal{M}$ is well formed.

Proof To verify that the iDt has the structure of a simplicial complex we need to first ensure that each triangle is a simplex: it must have three distinct vertices. This amounts
to eliminating the possibility of loop edges. We must further verify that the intersection of any two simplices in the iDt is a simplex. In particular, the intersection of two edges must be either empty, or a single point. This corresponds to ensuring that no two edges share the same pair of endpoints. We must also convince ourselves that the intersection between any two triangles is either a single edge, a single vertex, or empty. But this is already ensured if we have eliminated loop edges and double edges: the triangles cannot share three edges because by assumption $P$ contains at least four points and $\mathcal{M}$ is a single component.

There are thus two violations of a proper triangulation to consider:

1. loop edges, and
2. multiple edges between two vertices.

By showing that these correspond to the two violations of a well formed Voronoi diagram we establish that a proper iDt is equivalent to a well formed iVd.

That loop edges in the iDt are equivalent to Voronoi cells that are not topological disks is established by Lemma 3.2. Also, Lemmas 3.4 and 3.5 establish a one to one correspondence between Voronoi edges and their dual Delaunay edges. Thus multiple edges between two vertices in the iDt is equivalent to the corresponding Voronoi cells sharing multiple Voronoi edges.

This theorem allows us to use Voronoi diagrams to characterize the meshes that will admit a proper Delaunay triangulation. The dual picture does not yield additional information, but it gives another way to view the situation. In particular the Voronoi perspective is a natural one to have in mind when formulating sampling criteria.

### 3.2 Self-Delaunay meshes

Recognizing the nice properties of iDts on pwf surfaces as identified and motivated by several recent works [BS07, Gli05, DHLM05, WMKG07], and discussed in Section 3.1, we are naturally lead to the problem of exploiting such triangulations in practice. In the practical setting the pwf surface is invariably given in the form of a mesh, $M$.

An algorithm and data structure to generate the iDt of a pwf surface was presented by Fisher et al. [FSBS06]. The disadvantage of this approach is that a separate, somewhat complicated, data structure must be maintained to describe the iDt, in addition to the
extrinsic triangulation inherent in the mesh itself. The obvious solution to this problem is to seek a mesh representation $M$ such that the iDt of $M$ coincides with its extrinsic triangulation. Such a mesh is a self-Delaunay mesh, which we will occasionally abbreviate as sDm.

We have seen that any pwf surface, $\mathcal{M}$, admits an iDt, but in general the triangulation is not proper. We require that the vertex set $P$ of the iDt include all of the cone points of $\mathcal{M}$, however we allow $P$ to also contain points that are not cone points. In that sense the iDt of $P$ on $\mathcal{M}$ is not entirely intrinsically defined: it requires the specification of $P$.

A manifold Delaunay complex is a manifold simplicial complex that is a Delaunay complex with respect to its intrinsic metric. In particular, a pwf surface triangulated with its iDt is a manifold Delaunay complex if its iDt is proper. For any $P \subset \mathbb{R}^{d}, \mathcal{D}^{d}(P)$ is a manifold Delaunay complex, but as we will see in Section 3.3, a manifold subcomplex of $\mathcal{D}^{d}(P)$ is not a manifold Delaunay complex in general.

Definition 3.9 A self-Delaunay mesh is a mesh $M$ whose associated simplicial complex $\mathcal{K}_{M}$ is a manifold Delaunay complex.

Thus $\mathcal{K}_{M}$, which defines the extrinsic triangulation of $M$, is also an iDt of the underlying pwf surface. We can express the definition of a self-Delaunay mesh in any of the equivalent ways that the iDt of $\mathcal{M}$ is defined. A self-Delaunay mesh is a mesh for which each triangle face has an empty geodesic circumdisk. A self-Delaunay mesh is a mesh in which each edge connects vertices that are neighbours in the iVd. It follows that the iVd of a self-Delaunay mesh is always well formed.

The characterization of a self-Delaunay mesh that is probably most algorithmically convenient is given by the angle sum property. A self-Delaunay mesh is a mesh for which each edge subtends a pair of angles whose sum does not exceed $\pi$.

Unless stated otherwise, $M$ is assumed to be without boundary, however the definition of a self-Delaunay mesh may be extended to include meshes with boundaries. One alternative is to demand the locally Delaunay property only for interior edges [BS07]. We make a stronger requirement. A manifold simplicial complex with boundary, $\mathcal{K}_{M}$, is a manifold Delaunay complex with boundary if every internal edge is locally Delaunay and every boundary edge is subtended by an angle not exceeding $\pi / 2$. A self-Delaunay mesh with boundary is one with such an associated simplicial complex. This definition ensures that boundary vertices are neighbours only if they are neighbours in the iVd. It also maintains the positive edge
weight property of the cotan operator under the assumption of Von Neumann boundary conditions [VL08].

Although the self-Delaunay mesh is defined in terms of its own intrinsic metric and is independent of a reference surface, it is not intrinsically defined. Indeed, no mesh can be intrinsically defined because the definition of a mesh involves a piecewise linear mapping into $\mathbb{R}^{3}$ via the extrinsic triangulation defined by the associated manifold simplicial complex, $\mathcal{K}_{M}$. The definition of a self-Delaunay mesh demands a marriage of intrinsic and extrinsic triangulations: the extrinsic triangulation defined by $\mathcal{K}_{M}$ is required to agree with the iDt of the underlying pwf surface.

### 3.2.1 Delaunay extrinsic edge flips

The edge flipping algorithm which produces an iDt on a pwf surface has been well studied [ILTC01, BS07, Gli05, FSBS06]. In this algorithm the triangulation of $\mathcal{M}$ is altered, but the geometry of $\mathcal{M}$ remains fixed, and the situation is similar to that of the traditional planar edge flipping algorithm [Law77].

We now consider a variation on this algorithm wherein we flip a mesh edge if it is nlD. Such an edge flip results in a change in the geometry of the mesh. The new mesh has a different underlying pwf surface. As a consequence the theoretical properties of this algorithm are very different from the case of fixed geometry considered earlier. We call such a mesh edge flip a Delaunay extrinsic edge flip.

The termination properties of this algorithm, and its employment in producing selfDelaunay meshes will be discussed in Chapter 6. We introduce the basic ideas and insights behind this simple algorithm here because it is a convenient theoretical tool for examining the properties of self-Delaunay meshes. In particular it will become evident that a selfDelaunay mesh on vertex set $P \subset S$ is not unique in general, even given the requirement that it be homeomorphic to $S$. Also, the edge flipping algorithm will play an important role in Chapter 4 when we consider the relationship between self-Delaunay meshes and Gabriel meshes.

The hinge defined by an edge $e=[p, q] \subset M$ is $e$ together with the two triangle faces adjacent to it. Edge $e$ is the pivot of the hinge. The four vertices of a hinge define a tetrahedron that we call the flip-tet. Let a hinge on $e$ be defined by triangles $t_{1}=[u, p, q]$ and $t_{2}=[v, q, p]$. We will refer to this hinge as a triple $(u, e, v)$, or a pair $\left(t_{1}, t_{2}\right)$, as convenient. Let $\sigma$ be the flip-tet defined by this hinge. The edge $e^{\prime}=[u, v]$ is the opposing
edge to $e$ and it defines, together with its adjacent faces in $\sigma$, an opposing hinge to that of $e$. In an extrinsic edge flip the flip-tet takes the place of the flip-quad employed in the fixed geometry case. In an extrinsic edge flip we replace hinge ( $u, e, v$ ) with its opposing hinge ( $p, e^{\prime}, q$ ).

Flipping an nlD edge always yields a locally Delaunay edge. This is easy to see since the sum of the angles subtending $e$ and $e^{\prime}$ is at most $2 \pi$. Referring to Figure 3.3, we have

$$
\begin{equation*}
\angle p u q+\angle u q v+\angle q v p+\angle v p u \leq 2 \pi, \tag{3.1}
\end{equation*}
$$

with equality holding only when $p, q, u$, and $v$ are coplanar.

Lemma 3.10 If edge $e$ in $M$ is not locally Delaunay, then its opposing edge $e^{\prime}$ is.


While Lemma 3.10 is true in this setting, its converse, which holds for any planar quad in general position [dBvKOS98], is not true in a mesh. For example, consider a hinge on an edge $e$ adjacent to two equilateral triangles such that the associated flip-tet is a the regular tetrahedron: both $e$ and its opposing edge are locally Delaunay. Consequently, there can be multiple self-Delaunay meshes on the same vertex set and defining the same topological surface. Thus in this sense, without demanding further qualifications, we do not have a general uniqueness theorem for self-Delaunay meshes, contrary to the case of fixed geometry, be it planar 2D or a fixed piecewise flat surface.

Although Lemma 3.10 ensures that any flippable edge in our algorithm would improve matters locally, it does not lead to a termination proof. This issue is addressed in Section 6.2.

### 3.2.2 Smooth self-Delaunay meshes

In contrast to the rDt and the iDt-mesh, self-Delaunay meshes do not depend upon $S$ for their definition. This is an attractive property, but it comes at a price. Because a selfDelaunay mesh is independently defined, it need not be a good representation of $S$, even if its vertices have been well sampled from $S$. We can connect the sample points in a way
that does not make sense for representing $S$, but till yields a self-Delaunay mesh. Figure 3.4 gives an example of such a situation.


Figure 3.4: Without some kind of regularity condition, a self-Delaunay mesh is not necessarily a good representation of the surface from which its vertices are sampled. The mesh fragment shown here is necessarily self-Delaunay since it has no obtuse face angles. However if the samples represent a well sampled smooth surface, the vertex $a$ should not be associated with the neighbours that it has. For example, $a$ could be a point on the antipodal side of a well sampled sphere.

Therefore we wish to focus our attention on a subclass of self-Delaunay meshes. As we have seen in Section 2.4, a mesh that is a good representation of $S$ must be a smooth mesh. Thus we are interested in smooth self-Delaunay meshes. While any self-Delaunay mesh that is a good representation of $S$ must be smooth, we do not have a demonstration of the converse, i.e., that a smooth self-Delaunay mesh must be a good representation of $S$. We believe that, given topological equivalence, any smooth mesh which interpolates all the vertices of $S$ must closely approximate $S$ when $P$ is a sufficiently good sample set, but this conjecture is left for future work.

### 3.3 Distinctions between the Delaunay structures

In this section we give specific examples that show that the rDt and the iDt-mesh are not self-Delaunay meshes in general. The example used to demonstrate the distinction between the iDt-mesh and a self-Delaunay mesh does not yield a distinction between the rDt and a self-Delaunay mesh. It follows immediately then that the rDt and the iDt-mesh are also distinct structures.

We proceed by constructing an nlD hinge that may appear in the mesh structure under investigation. The construction for the iDt-mesh is simpler than that required for the rDt, so we begin with that.

### 3.3.1 The iDt-mesh is not a self-Delaunay mesh

When we construct an iDt-mesh from a set $P$ of samples on $S$, then no matter what the density of $P$ on $S$, there will always be some metric distortion. Geodesic distances on the resulting iDt-mesh will not coincide exactly with geodesic distances between corresponding points on $S$. We exploit this to demonstrate that the iDt-mesh need not be a self-Delaunay mesh.


Figure 3.5: The cylinder example to illustrate the discrepancy between an iDt-mesh and a self-Delaunay mesh. (a) In a quadrilateral with opposing angles equal and all sides equal, the longer diagonal edge is nlD. When the quadrilateral is defined by linear interpolation between the vertices, this is edge $e^{\prime}=[u, v]$, but in its geodesic realization on the cylinder, the other diagonal, the (geodesic) circular arc $\{p, q\}$, is longer and therefore nlD. (b) A cross-sectional profile of the cylinder at geodesic edge $e=\{p, q\}$, where $s$ is the geodesic distance between $p$ and $q$.

To construct an example of this, consider a planar quadrilateral puqv such that all four sides are of equal length and the opposite angles are equal, i.e., $\angle u p v=\angle u q v$ and
$\angle p u q=\angle p v q$. Thus puqv is a diamond. Suppose further that one of the diagonals is slightly shorter than the other. Specifically, let $|e|=|[p, q]|=\ell$ and $\left|e^{\prime}\right|=|[u, v]|=\ell+\epsilon$. For the symmetric quad puqv, the longer diagonal edge $e^{\prime}$ is nlD since it is subtended by larger angles, i.e., $\angle u p v+\angle u q v>\angle p u q+\angle p v q$.

Consider a cylinder $S$ of radius $r$. Allow the quadrilateral to pivot on the diagonal $e^{\prime}$ and place its four vertices on the cylinder so that $e^{\prime}$ is parallel to the axis of the cylinder (Figure 3.5(a)). In the geodesic realization of the quadrilateral, the geodesic diagonal corresponding to $e$, drawn as the short circular arc between $p$ and $q$ in Figure 3.5(a), will have length $s=4 r \arcsin \left(\frac{\ell}{4 r}\right)$ (see Figure 3.5(b)). Thus its length will be longer than that of the other diagonal $e^{\prime}$ on the surface of $S$, where $\left|e^{\prime}\right|=\ell+\epsilon$, as long as $\frac{\ell}{4 r}>\sin \left(\frac{\ell+\epsilon}{4 r}\right)$. This is realizable for a sufficiently small $\epsilon$, even though sampling density requirements may constrain the size of $\frac{\ell}{4 r}$ to be small.

Indeed, as long as $\frac{\ell}{4 r}>0$, we can select $\epsilon$ such that $0<\frac{\epsilon}{4 r}<\arcsin \left(\frac{\ell}{4 r}\right)-\frac{\ell}{4 r}$. In practice, $\epsilon$ can be arbitrarily small while the samples would still technically be in general position. A small $\epsilon$ corresponds to Voronoi vertices that are very close together. In this case, $e^{\prime}$ is the locally Delaunay edge on the surface of the cylinder $S$ and present in the iDt of $S$. But it is nlD in the resulting iDt-mesh $M$, which would consequently not be a self-Delaunay mesh.

Note that a similar example could be constructed if $S$ were a pwf surface. In other words, if we were to take a given mesh $M$ and produce a new mesh $M^{\prime}$ with the same vertices, but with connectivity defined by the iDt of $M$, then $M^{\prime}$ would not be a self-Delaunay mesh in general, even if $M$ has a well formed Voronoi diagram.

### 3.3.2 The rDt is not a self-Delaunay mesh

Using similar arguments as for iDTs we can show that the restricted Delaunay triangulation is not a self-Delaunay mesh in general.

In order to construct an example where the restricted Delaunay triangulation yields an edge that is nlD in the resulting mesh, we again make use of the diamond-shaped quadrilateral puqv as defined in Section 3.3.1. Let the quadrilateral be bent at the nlD edge $e^{\prime}=[u, v]$ by an angle of $2 \alpha$ and inscribe the vertices $u, v, p, q$ on an ellipsoid $S$ such that $e^{\prime}$ is parallel to the principle axis of $S$, as shown in Figure 3.6(a). Let $\sigma$ be the flip-tet associated with $\left(p, e^{\prime}, q\right)$, and let $c_{\sigma}$ be its circumcentre. It is not hard to show directly that $c_{\sigma}$ lies outside of $\sigma$. (The assertion also follows immediately from Lemma 4.7 in Chapter 4.) Furthermore, $c_{\sigma}$ sits right above $c_{e^{\prime}}$, the midpoint of $e^{\prime}$; let it be at a distance $z$ away. By considering
triangles $\left[c_{\sigma}, c_{e^{\prime}}, u\right]$ and $\left[c_{\sigma}, c_{e^{\prime}}, q\right]$, a straight forward calculation yields $z=\left(\epsilon+\frac{\epsilon^{2}}{2 \ell}\right) /(2 \sin \alpha)$.


Figure 3.6: The ellipsoid example to illustrate the discrepancy between an rDt and a selfDelaunay mesh. (a) The geometry of the tetrahedron puqv inscribed in an ellipsoid $S$. (b) A cross-sectional profile of the ellipsoid perpendicular to edge $e^{\prime}$, where $c$ is the circumcentre of the tetrahedron and it lies above $e^{\prime}$ but inside the ellipsoid.

Although sampling density constraints, i.e., based on a sampling radius moderated by the local feature size, can force $\alpha>0$ to be small, we can still make $z>0$ arbitrarily small by choosing a sufficiently small $\epsilon$ (for fixed $\ell, \alpha$ is bounded away from zero). Since there is a gap between the ellipsoid surface $S$ and edge $e^{\prime}$, we can make $c_{\sigma}$ lie inside the ellipsoid, as shown in Figure 3.6(b). Consequently, $S$ would pass above the circumcentre $c_{\sigma}$.

Now consider the 3D Voronoi cells $V(u)$ and $V(v)$. The Voronoi face $f$ which separates them lies on a plane which is a perpendicular bisector of edge $e^{\prime}$. This face extends upwards to infinity and must pass through the circumcentre $c_{\sigma}$, as $c_{\sigma}$ is equi-distant to $u, v, p$ and $q$. Since $c_{\sigma}$ is inside the ellipsoid $S$, the surface of $S$ must intersect face $f$. It follows that the restricted Voronoi cells $\left.V\right|_{S}(u)$ and $\left.V\right|_{S}(v)$ are neighbours and the nlD edge $e^{\prime}=[u, v]$ would appear in the rDt, which is therefore not a self-Delaunay mesh.

### 3.4 Discussion

When given a mesh $M$, one can determine with a quick check of the angles whether or not it is a self-Delaunay mesh. In contrast, there is no way to tell if it is a rDt or an iDt-mesh, unless the reference surface $S$ is also supplied. In this sense self-Delaunay meshes may be compared to cocone meshes, which also do not rely on $S$ for their definition. In order to verify that a mesh is a cocone mesh, one must construct $\mathcal{D}^{3}(P)$, the Delaunay
tetrahedralization of the ambient space. Self-Delaunay meshes do not require this overhead for their verification, but as yet there is no guarantee that a self-Delaunay mesh, even a smooth one, is a good representation for $S$.

We introduced the rDt and the iDt-mesh in Section 2.3 by considering them to be the results of employing different metrics on $S$ when applying the Delaunay paradigm. New Delaunay structures for $S$ may be created by employing new metrics in the definition. Many Delaunay-based surface representation cannot naturally be viewed as Delaunay structures in this sense. For example, cocone meshes do not easily lend themselves to this kind of description. However, a self-Delaunay mesh that is a good representation of $S$, can be viewed as a Delaunay structure for $S$ in this way.

Define a conforming homeomorphism as a homeomorphism $h: M \rightarrow S$ that leaves the samples in $P$ invariant: $\left.h\right|_{P}=\operatorname{id}_{\mathbb{R}^{3}}$. The canonical example of such a homeomorphism is $\xi_{S}$, the orthogonal projection onto $S$. If $h$ is a conforming homeomorphism for the selfDelaunay mesh $M$, then we can view the intrinsic metric on $M$ as a metric on $S$ : for $x, y \in S, d(x, y)=d_{M}\left(h^{-1}(x), h^{-1}(y)\right)$. We can then view $M$ as the dual mesh to the Voronoi diagram of its vertices on $S$ with respect to this metric. Such a characterization of a self-Delaunay mesh is somewhat contrived, but it provides a convenient framework for our discussion.

Guibas and Russel [GR04] demonstrated that the planar Delaunay triangulation can be surprisingly brittle under small perturbations of the sample set. In other words a small change in the input point set can yield a large change in the connectivity of the resulting Delaunay triangulation. The distinctions between the Delaunay structures that we illuminated in Section 3.3 can be viewed as a manifestation of this same phenomenon. Whereas Guibas and Russel considered perturbations of the vertices, we considered perturbations of the metric.


Figure 3.7: Voronoi distortion

The issue can be seen as arising from samples that are "almost not in general position". We can characterize this situation as occurring when Voronoi vertices become very close. Although the assumption that the samples are in general position with respect to a given metric may be justified (see the discussion in Section C.1), we cannot assume any predefined minimum distance between Voronoi vertices.

All three Delaunay structures, with the appropriate qualifications, are expected to provide good representations of $S$. As the sampling density increases, the differences between the associated metrics over the span of a Voronoi cell will decrease. However, there will always be small differences and the metric perturbation near two Voronoi vertices that are very close in one metric may result in a Voronoi edge disappearing in another metric, as schematically depicted in Figure 3.7. Such discrepancies are expected to arise in the neighbourhood of sliver tetrahedra in $\mathcal{D}^{3}(P)$ for example.

### 3.4.1 A fourth Delaunay surface structure

| metric | on $S$ | on $M$ |
| :--- | :---: | :---: |
| extrinsic | rDt | self-rDt |
| intrinsic | iDt-mesh | sDm |

Table 3.1: A family of Delaunay structures

We have considered the iDt-mesh, the rDt, and the self-Delaunay mesh, Delaunay structures arising on $S$ as a result of using $d_{S}, d_{\mathbb{R}^{3}} \mid S \times S$, and $d_{M}$ respectively. In other words we have considered the intrinsic metric on $S$, the extrinsic metric on $S$, and the intrinsic metric on $M$. In this context it is natural to consider also the extrinsic metric on $M$, i.e., $\left.d_{\mathbb{R}^{3}}\right|_{M \times M}$. This would correspond to measuring the distance between points on $M$ with respect to $d_{\mathbb{R}^{3}}$. We may call a Delaunay structure arising from this metric a self-rDt mesh. Indeed, such a mesh would be a restricted Delaunay triangulation of itself. This family of four Delaunay structures is summarized in Table 3.1.

We don't pursue a study of the self-rDt mesh in this thesis, although we will point out in Section 4.7 that there is no "locally self-rDt" property of a mesh edge that can be decided on the basis of its hinge alone. This means that there can be no simple extrinsic edge flip algorithm to find self-rDt meshes. We are not aware of any explicit investigations into selfrDt meshes in the literature, however there are indications that these structures may have close ties to the flow complex.

## Chapter 4

## Gabriel meshes are self-Delaunay meshes

In this chapter we provide a comparison between the local properties of Gabriel meshes and self-Delaunay meshes. Recall that a Gabriel mesh is a triangle mesh each of whose faces has a diametric ball empty of mesh vertices (an empty diametric ball). We say the faces have the Gabriel property with respect to the vertex set.

The Gabriel property is defined as a global condition, but it is misleading to emphasize this if we are dealing with points well sampled from a smooth surface. In this case the circumradii involved are a small fraction of the distance to the medial axis: geodesically distant points are irrelevant. What is at issue is the local connectivity of the samples.

Both the Gabriel mesh and the self-Delaunay mesh can be seen as attempts to extend the Delaunay paradigm to manifold meshes. Consider a set of sample points lying in a plane in $\mathbb{R}^{3}$. A Gabriel mesh on these samples is equivalent to a Delaunay triangulation of the points because the restriction of the diametric ball of triangle $t$ to the plane is just the circumdisk of $t$. If the circumdisks are empty then so will be the diametric balls, and vice versa. Likewise, the circumdisks of the triangles are geodesic disks on the planar mesh, so for planar meshes, there is no distinction between a Gabriel mesh and a self-Delaunay mesh. The distinction between self-Delaunay meshes and Gabriel meshes arises when we consider general manifold meshes.

This study of Gabriel meshes was inspired by the work of Cheng and Dey [CD07], where they describe an edge flipping algorithm that seeks to produce a Gabriel mesh from a given
mesh whose vertices are well sampled from a smooth surface. We provide an analysis of this algorithm that reveals a close relationship with the Delaunay edge flipping algorithm introduced in Section 3.2.1. The local criteria that are required of a Gabriel mesh are slightly stronger than what is required of a self-Delaunay mesh, and this is important, because a closed Gabriel mesh does not exist in general [Cha03]. In fact, we show in Section 4.5 that even when local uniformity constraints are imposed on the sample set, obstructions to the existence of closed Gabriel meshes may remain.

The authors of [CD07], which is a technical report, have retracted their result, but the motivation behind that work remains. The appeal of the Gabriel mesh is that it is a locally defined surface representation: The connectivity of the vertices is constrained simply in terms of essentially local information and without the need for an independent reference surface. In particular it is not necessary to construct the full Delaunay triangulation of the ambient three dimensional space in order to verify the local connectivity. As we have discussed in Chapter 3, this characteristic is also shared by self-Delaunay meshes.

We show that Gabriel meshes are self-Delaunay meshes. In Section 4.4 we demonstrate that the locally Delaunay property of an edge is a slight relaxation of what is locally demanded by a Gabriel mesh. We also introduce, in Section 4.2, a definition of the Gabriel complexes that reveals a natural hierarchical decomposition of the Delaunay triangulation, which we express in Theorem 4.1. This description of the Gabriel complexes suggests they may hold promise as a scaffold for a manifold reconstruction algorithm in a high dimensional ambient space, where the full Delaunay triangulation is impractical.

### 4.1 Background

Unlike self-Delaunay meshes, Gabriel meshes have seen some attention in the context of surface reconstruction. Their simple definition ensures that they will be substructures of $\mathcal{D}^{3}(P)$, which in turn implies that smooth Gabriel meshes are likely to be cocone meshes, and so benefit from the approximation guarantees they enjoy.

Petitjean and Boyer [PB01] defined the Gabriel complex of a given set of points $P$ as consisting of those triangles in $\mathbb{R}^{3}$ that have the Gabriel property with respect to $P$. A reconstruction algorithm was introduced which extracted a manifold triangle mesh from the Gabriel complex. A heuristic argument, based on the planar case mentioned above, was given for why there should be enough triangles in the Gabriel complex to extract a closed
manifold mesh if $P$ is well sampled from a smooth surface. The issue of near degenerate configurations was not mentioned, although it was recognized to be a problem with an algorithm described earlier in the paper.

In their surface reconstruction algorithm Adamy et al. [AGJ00] also implicitly assumed that it was possible to extract full umbrellas of Gabriel triangles around each vertex. A topological clean-up step filled holes in the extracted mesh, some of which were produced by other topological reparation steps. The Gabriel complex was defined as the simplices of codimension one that had the Gabriel property and ambient dimensions two and three were considered. Gumhold et al. [GWM01] subsequently also employed the Gabriel-based methods introduced in [AGJ00].

The surface reconstruction algorithm presented by Attene and Spagnuolo [AS00] also exploited the Gabriel property. Interestingly it was not assumed that the Gabriel faces would form a closed surface, but that holes would be limited to isolated missing triangles.

Then Chaine [Cha03] observed that a certain sliver tetrahedron in the 3D Delaunay tetrahedralization was sufficient to prevent full umbrellas of Gabriel faces at its vertices. We refer to such a tetrahedron as a tetrahedral obstruction to Gabriel meshes.

### 4.2 Gabriel complexes

We begin by defining the Gabriel complexes in arbitrary dimensions. Let $P \subset \mathbb{R}^{d}$ be a finite set of sample points in general position. Let $k \leq d$ and let $\sigma_{k}$ be a $k$-dimensional simplex with vertices in $P$. The diametric $d$-ball for $\sigma_{k}$, denoted $B_{\sigma_{k}}$, is the smallest d-dimensional ball containing the vertices of $\sigma_{k}$ on its boundary: $B_{\sigma_{k}}=B_{\mathbb{R}^{d}}\left(c_{\sigma_{k}} ; r_{\sigma_{k}}\right)$, where the centre, $c_{\sigma_{k}}$, and the radius, $r_{\sigma_{k}}$, are the circumcentre and circumradius of $\sigma_{k}$.

We say $\sigma_{k}$ is $k$-Gabriel if its diametric $d$-ball does not contain any points of $P$. The $k$ Gabriel complex of $P$ is the simplicial complex formed by the $k$-Gabriel simplices and their faces. Thus the 1-Gabriel complex is the Gabriel graph of $P$, and the $d$-Gabriel complex is the Delaunay triangulation of $P$. For completeness, define the 0 -Gabriel complex to be the sample points themselves. Denote by $\mathcal{G}^{k}(P)$ the $k$-Gabriel complex. Then we have

Theorem 4.1 The Gabriel complexes form a nested hierarchy of subcomplexes of the Delaunay triangulation of $P \subset \mathbb{R}^{d}$ :

$$
P=\mathcal{G}^{0}(P) \subset \mathcal{G}^{1}(P) \subset \cdots \subset \mathcal{G}^{k}(P) \subset \cdots \subset \mathcal{G}^{d}(P)=\mathcal{D}^{d}(P) .
$$

That $\mathcal{G}^{0}(P) \subset \mathcal{G}^{1}(P)$ can be established by demonstrating that the Gabriel graph contains the nearest neighbour graph. Our proof of Theorem 4.1 is an extension of this method. The edge to the nearest neighbour of a point $p$ can be characterized as the 1 -simplex on $p$ that has the smallest diametric $d$-ball. This ball is necessarily empty of sample points. The extension to higher dimensions follows by showing that if $\sigma_{k}$ is $k$-Gabriel, then the $(k+1)$ simplex with the smallest diametric $d$-ball amongst those that have $\sigma_{k}$ as a face, must be $(k+1)$-Gabriel. The technical demonstration of this result is postponed to Section 4.6.

According to Theorem 4.1, the $k$-Gabriel complex is a subcomplex of the $k^{\prime}$-Gabriel complex if $k^{\prime} \geq k$. However not all the $k$-simplices of the $k^{\prime}$-Gabriel complex need belong to the $k$-Gabriel complex. It is this latter fact that motivates the introduction of the $k$ prefix.

Indeed, we will be interested in the edges of a substructure of the 2-Gabriel complex, but these edges need not belong to the Gabriel graph. Triangle faces containing an edge $e$ of the 1-Gabriel complex cannot have an obtuse angle subtended by $e$, but the 2-Gabriel complex may well contain obtuse triangles. This property of 1-Gabriel edges implies that they will never be flipped by a Delaunay edge flip, and so we anticipate their relevance in this context. However, for the purposes of this chapter, an unspecified mention of the Gabriel property can be understood as a reference to the 2-Gabriel complex.

In the rest of this chapter, until Section 4.6 , we will take $d=3$ for convenience. Note, however, that the ambient dimension has no bearing on the results. Local computations involving an isolated flip-tet may be confined to the affine hull of the tetrahedron.

Definition 4.2 (Gabriel mesh) A Gabriel mesh is a manifold triangle mesh that is a substructure of the 2-Gabriel complex.

Since Gabriel triangles have empty circumballs, they belong to the Delaunay tetrahedralization by definition. We have the following useful characterization:

Lemma 4.3 A triangle $t \subset \mathbb{R}^{3}$ is Gabriel iff its dual Voronoi edge intersects the affine hull of $t$.

Proof If $t$ is Gabriel, the diametric $d$-ball of $t$ is empty and so its centre must lie on the Voronoi edge dual to $t$. But the centre of the diametric ball also lies on the plane defined by $t$. Conversely, if the Voronoi edge dual to $t$ intersects the affine hull of $t$ at $c_{t}$, then the empty circumball of $t$ centred at $c_{t}$ is a diametric ball, and so $t$ is Gabriel.

### 4.2.1 Gabriel faces in a tetrahedron

In the next section we will prepare for a detailed examination of hinges. Every hinge has an associated flip-tet, and we make some preliminary observations here about the Gabriel properties of tetrahedra. We will assume that the tetrahedron is non-degenerate. The case where the affine hull of a flip-tet is a plane corresponds to the usual case of planar Delaunay edge flips. More extreme degeneracies can be dealt with by the same arguments that permit the assumption of general position for planar point sets.

Let $\sigma$ be a tetrahedron with circumsphere $S_{\sigma}$ and circumcentre $c_{\sigma}$. The interior halfspace with respect to $\sigma$ of a triangle face $t \subset \sigma$ is the half space bounded by the supporting plane of $t$ and containing the fourth vertex of $\sigma$. Likewise, the exterior half-space of $t$ is the one that does not contain the fourth vertex. We define these half spaces to be closed: the intersection of the interior and the exterior half spaces of $t$ is the affine hull of $t$.

The following two lemmas concern the Gabriel properties of the faces of an isolated tetrahedron. In other words we consider the sample set to be comprised solely of the four vertices of the tetrahedron.

Lemma 4.4 A face $t$ in a tetrahedron, $\sigma$, is Gabriel iff $c_{\sigma}$ lies in its interior half-space.

Proof Consider the Voronoi diagram of the vertices of $\sigma$. A Voronoi edge is supported by the line perpendicular to its dual face and through its circumcentre. The Voronoi edges all extend to infinity in one direction and terminate at $c_{\sigma}$ in the other. A Voronoi edge dual to $t$ must extend to infinity in the exterior half-space of $t$ : given a sufficiently large radius a ball centred in this half-space and with the vertices of $t$ on its boundary will not contain the fourth vertex of $\sigma$. Thus the Voronoi edge dual to $t$ terminates at $c_{\sigma}$ and always extends to the exterior half space of $t$. Therefore it will intersect the plane supporting $t$ iff $c_{\sigma}$ lies in interior half-space of $t$. The result follows from Lemma 4.3.

An alternate proof of Lemma 4.4 is given by Lemma 4.13, its generalization to higher dimensions. Lemma 4.4 is one of the two principal observational tools we use to extract our results. This characterization of Gabriel faces in a tetrahedron facilitates an understanding of the four point configurations that obstruct Gabriel meshes.

Lemma 4.5 Every tetrahedron has at least two Gabriel faces.

Proof Suppose $\sigma$ does not have two Gabriel faces. Lemma 4.4 implies that $c_{\sigma}$ lies in the intersection of the exterior half spaces of three faces of $\sigma$. Let $v$ be the vertex common to the three faces. Then $c_{\sigma}$ is contained within a space exterior to $\sigma$ and defined by a solid angle opposite to (thus congruent to) the solid angle of $\sigma$ at $v$. Since the tetrahedron is convex, this solid angle is less than $2 \pi$ and it follows that a vector from $c_{\sigma}$ to $v$ would point towards the interior of $\sigma$. But the direction of the exterior normal vector of $S_{\sigma}$ at $v$ coincides with that of the vector from $c_{\sigma}$ to $v$, and cannot point towards the interior of $\sigma$. Thus we have a contradiction.

### 4.3 The anatomy of a hinge

We will examine the distinction between Gabriel meshes and self-Delaunay meshes by focusing on edges and their adjacent faces: hinges. A mesh is not a self-Delaunay mesh if any of its edges is not locally Delaunay. We will compare this local Delaunay criterion with the natural corresponding condition that edges in a Gabriel mesh must possess. We say that $e$ obtains a Gabriel certificate from each of its adjacent faces that is Gabriel in the associated flip-tet. Edge $e$ will be locally Gabriel only if it has two Gabriel certificates.

In all that follows, the hinge will consist of edge $e=[p, q]$ and adjacent triangles $t_{1}=$ $[p, q, u]$ and $t_{2}=[q, p, v]$. It is convenient to work with spheres rather than balls. Denote by $S_{t}$ boundary of a diametric ball $B_{t}$. The flip-tet associated with the hinge is $\sigma$. Its circumsphere is $S_{\sigma}$ with centre $c_{\sigma}$.

Faces $t_{1}$ and $t_{2}$ have consistent normals if they are both oriented towards their respective interior half-spaces, or both towards their exterior half-spaces with respect to $\sigma$. The dihedral angle of a hinge is the angle between consistent normals to the faces ${ }^{1}$. A hinge is non-sharp if the dihedral angle is less than $\pi / 2$. In all that follows we consider only non-sharp hinges. Since our interest is in smooth meshes, this is not a limitation.

### 4.3.1 Equivalent hinges

To test the locally Delaunay property of $e$ it is sufficient to check the sum of the angles it subtends. Two different hinges on $e$ will yield the same result on the Delaunay test if their

[^6]angles are the same. This reflects the fact that what is important is the circumcircles of the triangles, not the triangles themselves.

For the Gabriel property, not only are the angles subtended by $e$ important, but also the dihedral angle between $t_{1}$ and $t_{2}$. Given two circumcircles on pivot $e$, the dihedral angle will affect $S_{\sigma}$ and the relative position of $c_{\sigma}$. We say two hinges, $\left(t_{1}, t_{2}\right)$ and $\left(t_{1}^{\prime}, t_{2}^{\prime}\right)$ on $e$ are equivalent if the dihedral angle between $t_{1}$ and $t_{2}$ is the same as that between $t_{1}^{\prime}$ and $t_{2}^{\prime}$ and the angle subtended by $e$ in $t_{i}$ is equal to that in $t_{i}^{\prime}, i=1,2$.

Thus equivalent hinges have the same $S_{\sigma}$, and the affine hulls of their faces coincide. The circumcircles (intersections of


Figure 4.1: Equivalent hinges the affine hulls with $S_{\sigma}$ ) are the same, as are the interior and exterior half-spaces. Equivalent hinges yield the same Gabriel certificates and have the same locally Delaunay status. Note that this does not imply equivalent consequences if we were to perform an edge flip. We will not be concerned with the flip-tet itself until Section 4.5.

### 4.3.2 Cross-sectional diagrams

Having established that we only need to consider equivalence classes of hinges, we choose the most convenient representative from each class for the purposes of visualization and analysis. A hinge on $e$ is the canonical representative of its equivalence class if $u$ and $v$ lie on the perpendicular bisector plane of $e$.

This bisector plane defines a cross-section of the hinge and its circumsphere. A crosssectional diagram is our visualization of these cross sections, which we now describe. We caution that some objects in a cross-sectional diagram are one dimensional representations of their higher dimensional counterparts in the hinge, yet we use the same symbol for both. Thus, for example, $t_{1}$ may be a segment in the diagram or a triangle in the hinge: we rely on context to resolve the ambiguity.

Refer to Figure 4.2. The circumsphere of the hinge, $S_{\sigma}$, is depicted as the large solid
circle. Triangle $t_{1}$ is depicted as a solid segment between $u$ and point $e$. The latter point is the centre of the pivot edge of the hinge. The circumcircle of $t_{1}$ is represented by its diameter, the segment $\left[u, u^{*}\right]$. The portion of this diameter that does not contain $t_{1}$ is drawn with a dashed line and denoted $t_{1}^{*}$. If $e$ subtends an acute angle in $t_{1}$, then $\left|t_{1}\right|>\left|t_{1}^{*}\right|$ in the diagram. The other triangle in the hinge, $t_{2}$, is represented in the same way.

Figure 4.2 is used as a visual aid to the proof of Lemma 4.8 which says that if $e$ has two Gabriel certificates, then it is locally Delaunay. Appealing to Lemma 4.4, we see that the hinge drawn in the figure has two Gabriel certificates because $c_{\sigma}$ lies in the intersection of the interior half-spaces of $t_{1}$ and $t_{2}$. Indeed, it follows directly from the definition that the interior half-space of $t_{1}$ is represented in our diagram by the half-plane bounded by the affine hull of $t_{1}$ and containing $t_{2}$.

The Delaunay condition of a hinge can be checked by unfolding the faces into a common plane and determining if the circumcircle of one triangle contains the other triangle. In terms of our diagram, Figure 4.2, we could rotate $t_{2}$ clockwise on $e$ until it lies in the affine hull of $t_{1}$ and then check if $v$ in this new position lies outside of $S_{\sigma}$. In other words $e$ will be locally Delaunay if $u^{*}$ lies inside the circle of radius $\left|t_{2}\right|$ centred at $e$. This is just a check on the relative lengths of $t_{2}$ and $t_{1}^{*}$. Thus, an equivalent check is to test whether $v$ itself lies outside the circle of radius $\left|t_{1}^{*}\right|$ centred at $e$. Since this is the case in Figure 4.2, the edge $e$ shown there is locally Delaunay.

This leads to another observation we can make from the diagrams. When we draw a hinge, $\left(t_{1}, t_{2}\right)$ on $e$, in this fashion we at the same time draw another hinge $\left(t_{1}^{*}, t_{2}^{*}\right)$ on $e$ represented with dotted lines. We call this hinge the complementary hinge. The dihedral angle of a hinge is the same as that of its complementary hinge. If $e$ subtends an acute angle in $t_{1}$, then it will subtend an obtuse angle in $t_{1}^{*}$ (it has the complementary angle subtended by $e$ ). Also, if $t_{1}$ yields a Gabriel certificate to $e$, then $t_{1}^{*}$ will not, and vice versa. If $e$ is locally Delaunay with respect to $\left(t_{1}, t_{2}\right)$, then it will be not locally Delaunay with respect to $\left(t_{1}^{*}, t_{2}^{*}\right)$, and vice versa. Thus results about the former hinge directly yield "complementary results" about the latter.

So, for example, the first part of Lemma 4.7 says that it is impossible to draw a non-sharp hinge with $t_{1}$ and $t_{2}$ acute without having $c_{\sigma}$ contained in the intersection of their interior half-spaces. This immediately gives us the complementary result that it is impossible to draw a non-sharp hinge with $t_{1}^{*}$ and $t_{2}^{*}$ obtuse and such that $c_{\sigma}$ is contained in either of their interior half-spaces.


Figure 4.2: Lemma 4.8: If $e$ is locally Gabriel, then it is locally Delaunay.

### 4.4 Relating Gabriel and Delaunay properties

We now examine the relationship between Gabriel certificates yielded to $e$ and the locally Delaunay property of $e$. We emphasise that our results here apply only to non-sharp edges. We find the results naturally separate into two cases. If $t_{1}$ and $t_{2}$ both agree on the decision of whether or not to yield a Gabriel certificate, then we call it a symmetric case. We find there is no distinction between the locally Delaunay condition and the locally Gabriel condition here. If the faces don't agree, then we have an asymmetric case, and it is within this realm that we find the local distinction between self-Delaunay meshes and Gabriel meshes.

We will make reference to the following observation:
Lemma 4.6 Let $\ell$ be a chord in circle $S_{\sigma}$ and $e$ a point on it. If $c_{\ell}$ is the centre of $\ell$, then $\angle c_{\ell} e c_{\sigma}$ is acute.

### 4.4.1 Symmetric cases

We observe that much can be said by simply examining the angles subtended by $e$ :

Lemma 4.7 If both angles subtended by $e$ are acute, then $e$ has two Gabriel certificates. If these angles are both obtuse, then $e$ has no Gabriel certificates.

Proof Assume $t_{1}$ and $t_{2}$ have acute angles subtended by $e$. By Lemma 4.6, $\left[c_{\sigma}, e\right]$ makes an acute angle with both $t_{1}$ and $t_{2}$. Since the hinge is non-sharp, $t_{1}$ and $t_{2}$ cannot lie on the same side of the line supporting $\left[c_{\sigma}, e\right]$. Thus $c_{\sigma}$ must lie in the intersection of the interior half-spaces of $t_{1}$ and $t_{2}$, and so by Lemma $4.4, e$ is locally Gabriel. The complementary result follows.

The intermediate case, where both subtended angles are $\pi / 2$, occurs when $S_{t_{1}}$ and $S_{t_{2}}$ coincide ( $c_{\sigma}$ lies on $e$ and $S_{\sigma}$ is the diametric ball of both triangles), and so $e$ will be locally Gabriel. Thus if the angles are both obtuse or both nonobtuse, locally Delaunay and locally Gabriel mean the same thing. But we can say more:

Lemma 4.8 If $e$ is locally Gabriel, then it is locally Delaunay. If $e$ has no Gabriel certificates, then it is not locally Delaunay.

Proof Assume $e$ is locally Gabriel. In light of Lemma 4.7, we only need to check the case where $e$ subtends an acute angle in $t_{1}$ and an obtuse angle in $t_{2}$. We refer to Figure 4.2. We will show that $v$ must lie outside of the circle with centre $e$ and radius $\left|t_{1}^{*}\right|$. This circle, call it $C$, intersects $S_{\sigma}$ at $u^{*}$. If $c_{\sigma}$ lies on $t_{1}$ then $C$ will be tangential to $S_{\sigma}$ at $u^{*}$ and thus lie entirely within $S_{\sigma}$ and $e$ will be locally Delaunay. So assume $c_{\sigma}$ lies in the interior of the interior half-space of $t_{1}$. In this case $C$ will intersect $S_{\sigma}$ transversely at $u^{*}$. Furthermore, in a neighbourhood of $u^{*}$ it will be inside $S_{\sigma}$ in the interior half-space of $t_{1}$. Since $\left|t_{1}^{*}\right|<\left|t_{1}\right|$ by hypothesis, $C$ must intersect $t_{1}$ and thus must remain inside $S_{\sigma}$ within the interior halfspace of $t_{1}$. Since $t_{2}$ lies within this interior half-space, $C$ must intersect $t_{2}$ and thus $e$ is locally Delaunay. The complementary result follows.

Every edge in a Gabriel mesh must necessarily be locally Gabriel. We reiterate the implication of Lemma 4.8 for emphasis:

Theorem 4.9 A Gabriel mesh is a self-Delaunay mesh.

### 4.4.2 Asymmetric cases

If $e$ has a solitary Gabriel certificate, it must come from a triangle with an obtuse angle subtended by $e$ :


Figure 4.3: Lemma 4.10: If a triangle presenting an acute edge to $e$ yields a Gabriel certificate, so too must the other triangle.

Lemma 4.10 Suppose $t_{1}$ has an acute angle subtended by $e$. If $t_{1}$ yields a Gabriel certificate to $e$, then so too must $t_{2}$.

Proof Assume $t_{1}$ yields a Gabriel certificate to $e$, so $c_{\sigma}$ lies in its interior half-space. We will show that $t_{2}$ must also yield a Gabriel certificate to $e$.

Suppose to the contrary that $c_{\sigma}$ lies in the exterior half-space of $t_{2}$, and refer to the diagram of Figure 4.3. Let $\ell$ be the segment composed of $t_{2}$ and $t_{2}^{*}$. Then $\ell$ must separate $c_{\sigma}$ and $t_{1}$, since by definition $t_{1}$ lies within the interior half-space of $t_{2}$. By Lemma 4.6, and our assumption on $t_{1},\left[c_{\sigma}, e\right]$ makes an acute angle with $t_{1}$. Since the hinge is non-sharp, it must be $t_{2}^{*}$ rather than $t_{2}$ that lies between $\left[c_{\sigma}, e\right]$ and $t_{1}$. But by definition $t_{2}^{*}$ lies in the exterior half-space of $t_{1}$, contradicting the hypothesis that $t_{1}$ yields a Gabriel certificate to $e$.

When $e$ has a single Gabriel certificate, it is not locally Gabriel, but may yet be locally Delaunay. The following observation characterizes the distinction.

Lemma 4.11 If edge $e$ subtends an obtuse angle in triangle $t_{1}$ and an acute angle in $t_{2}$, then $e$ is nlD if and only if $r_{t_{1}}>r_{t_{2}}$.

Proof Unfold the hinge so that it is planar. Let $C_{1}$ and $C_{2}$ be the circumcircles of $t_{1}$ and $t_{2}$. Then the centre of $C_{1}$ will be on the same side of $e$ as $t_{2}$. Edge $e$ is nlD iff $C_{1}$ contains $C_{2}$ on that side, and by Lemma B.3, this happens iff $C_{1}$ has a larger radius than $C_{2}$.

### 4.5 Obstructions to Gabriel meshes

In this section we show that given a point set $P$, which may be nicely sampled from a smooth surface, it is in general not possible to construct a Gabriel mesh whose vertex set is $P$. We proceed by constructing a flip-tet whose opposing non-sharp hinges each have only a single Gabriel certificate. This is done by placing the tetrahedron, $\sigma$, at the equatorial plane, but to the side in such a way that $c_{\sigma}$ is exterior to $\sigma$, but close to a long sharp edge that is almost a diameter of $S_{\sigma}$. The tetrahedron is described in Figure 4.4.

This obstruction cannot be avoided by imposing any reasonable uniformity constraints on the sample set. A uniformity constraint imposes a lower limit on the distance between neighbouring sample points and thus imposes a lower bound on the edge lengths of triangles in a mesh. This uniformity constraint is coupled with the sampling radius which puts an upper bound on the triangle circumradius in any Delaunay-based reconstruction. Thus a uniformity constraint serves to put an upper bound on the ratio of the circumradius to shortest edge of the triangles. See [Dey07] for details. Note, however, that the shortest edge in the tetrahedron in Figure 4.4 can be made to be arbitrarily close in length to the radius of $S_{\sigma}$. Also, the hinges can be made arbitrarily flat.

If a Gabriel mesh exists on a sample set $P$ it need not be unique (consider a flip-tet that contains its circumcentre; both hinges are locally Gabriel, so it may be possible to flip that edge and obtain a different Gabriel mesh). Therefore even if we construct a mesh containing this obstruction and in which all faces outside of the obstruction are Gabriel, it doesn't immediately imply that a Gabriel mesh cannot be constructed: there may be a different Gabriel mesh in which no triangle is composed solely of vertices from $\sigma$. However, a Gabriel mesh is also a substructure of the Delaunay tetrahedralization, and as such we can constrain the possibilities by creating a sample set, $P$, in which $\sigma$ is the only sliver tet in the Delaunay tetrahedralization. In this way we see that a closed Gabriel mesh need not


Figure 4.4: Looking down on an obstruction to a Gabriel mesh. The figure is projected onto the equatorial plane. The non-sharp hinge on edge $[p, q]$ is on top. Edge $[v, q]$ lies on the equatorial plane while $p$ lies just above it and $u$ lies just below it. Thus $c_{\sigma}$ is in the interior half-space of neither $[p, v, q]$ nor $[u, v, q]$. It follows that neither $[p, q]$, nor its opposing edge $[u, v]$ have two Gabriel certificates: neither edge is locally Gabriel.
exist. A figure showing a similar obstruction embedded in a mesh of nice triangles can be found in [Cha03].

### 4.6 Proof of Theorem 4.1

The theorem follows directly from the following
Lemma 4.12 If $k<d$ and $\sigma_{k}$ is $k$-Gabriel, then it is a face of a $(k+1)$-Gabriel simplex.
We will first establish technical lemmas 4.13 and 4.14 , which are essentially observations relating to sphere intersections. Let aff $\left(\sigma_{k}\right)$ denote the affine hull of $\sigma_{k}$. In a triangulation, we say $\sigma_{k^{\prime}}$ is incident to $\sigma_{k}$ if $\sigma_{k}$ is a face of $\sigma_{k^{\prime}}$.

Let $\sigma_{k}$ be a face of $\sigma_{k+1}$. Then aff $\left(\sigma_{k}\right)$ defines two half-spaces in $\operatorname{aff}\left(\sigma_{k+1}\right)$. The interior half-space of $\sigma_{k}$ with respect to $\sigma_{k+1}$ is the one that contains the vertex of $\sigma_{k+1}$ that does
not belong to $\sigma_{k}$. We call the other half-space the exterior half-space. Contrary to balls and disks, we find it convenient to employ closed half-spaces: the interior and the exterior half spaces contain their boundaries.

This terminology provides a convenient local constraint on


Figure 4.5: Lemma 4.13 the $k$-Gabriel simplices:

Lemma 4.13 Let $\sigma_{k} \subset \sigma_{k+1}$, and let $v$ be the vertex of $\sigma_{k+1}$ that is not in $\sigma_{k}$. Then $v$ lies outside of $B_{\sigma_{k}}$ iff the circumcentre of $B_{\sigma_{k+1}}$ lies in the interior half-space of $\sigma_{k}$.

Proof This follows from Lemma B.4. Here $S_{1}$ and $S_{2}$ are the boundaries of $B_{\sigma_{k+1}}$ and $B_{\sigma_{k}}$ respectively, restricted to $\operatorname{aff}\left(\sigma_{k+1}\right)$. Since $v$ lies on $S_{1}$, it is exterior to $S_{2}$ if and only if it lies on the same side as the centre of $S_{1}$.

We need another observation that stems from standard sphere intersection properties:

Lemma 4.14 Suppose $\sigma_{k+1}$ and $\sigma_{k+1}^{\prime}$ share a face, $\sigma_{k}$, and that $c_{\sigma_{k+1}^{\prime}}$ is contained in the interior half-space of $\sigma_{k}$ with respect to $\sigma_{k+1}^{\prime}$. If $\sigma_{k+1}^{\prime}$ is contained in $B_{\sigma_{k+1}}$, then $r_{\sigma_{k+1}}>r_{\sigma_{k+1}^{\prime}}$.

Proof This is a consequence of Lemma B.5. Consider $S_{1}$ and $S_{2}$ to be the restriction of the boundaries of $B_{\sigma_{k+1}}$ and $B_{\sigma_{k+1}^{\prime}}$ respectively to aff $\left(\sigma_{k+1}^{\prime}\right)$. Let $v$ be the vertex of $\sigma_{k+1}^{\prime}$ that is not in $\sigma_{k}$. Since $\sigma_{k+1}^{\prime} \subset \bar{B}_{\sigma_{k+1}}$, it follows that $S_{2}$ is contained in $S_{1}$ on the side of $\operatorname{aff}\left(S_{1} \cap S_{2}\right)=\operatorname{aff}\left(\sigma_{k}\right)$ that contains $v$, which by hypothesis is also the side that contains the centre of $S_{2}$. Thus $S_{2}$ is smaller than $S_{1}$ and since its radius is exactly $r_{\sigma_{k+1}^{\prime}}$, it follows that $r_{\sigma_{k+1}}>r_{\sigma_{k+1}^{\prime}}$.

Finally, we demonstrate Lemma 4.12 by considering the simplex with the smallest diametric $d$-ball amongst those that


Figure 4.6: Lemma 4.14 have $\sigma_{k}$ as a face.

Proof of Lemma 4.12 By the definition of the Delaunay triangulation, $\sigma_{k}$ is a Delaunay $k$-simplex. Of the Delaunay $(k+1)$-simplices incident to $\sigma_{k}$, let $\sigma_{k+1}$ be the one with
the smallest diametric $d$-ball. We will show by induction that the vertices of all Delaunay simplices incident to $\sigma_{k+1}$ lie outside $B_{\sigma_{k+1}}$.

Since $\sigma_{k}$ is $k$-Gabriel, by Lemma 4.13, any $(k+1)$-simplex, $\sigma_{k+1}^{\prime}$, incident to $\sigma_{k}$ has its circumcentre in the interior half space of $\sigma_{k}$ with respect to $\sigma_{k+1}^{\prime}$. It follows then from Lemma 4.14 that the vertex of $\sigma_{k+1}^{\prime}$ that is not in $\sigma_{k}$ lies outside of $B_{\sigma_{k+1}}$.

Now consider the $(k+2)$-simplices incident to $\sigma_{k+1}$. These simplices must all have their vertices exterior to $B_{\sigma_{k+1}}$ because the vertex that is not in $\sigma_{k+1}$ belongs to another $(k+1)$-simplex that is incident to $\sigma_{k}$.

Let $k+2 \leq j<d$ and suppose all $j$-simplices incident to $\sigma_{k+1}$ have their vertices exterior to $B_{\sigma_{k+1}}$. Since every vertex in a $(j+1)$-simplex, $\sigma_{j+1}$, that is incident to $\sigma_{k+1}$ will also belong to a $j$-simplex that is incident to $\sigma_{k+1}$, it follows that all the vertices of $\sigma_{j+1}$ lie outside of $B_{\sigma_{k+1}}$.

Thus all Delaunay simplices incident to $\sigma_{k+1}$ have their vertices outside of $B_{\sigma_{k+1}}$. If $k+1=d$, then $\sigma_{k+1}$ is automatically $(k+1)$-Gabriel, because it is a Delaunay simplex. If $k+1<d$, then Lemma 4.15 below ensures that $\sigma_{k+1}$ is $(k+1)$-Gabriel.

Lemma 4.15 If $k<d$ and $\sigma_{k} \in \mathcal{D}^{d}(P)$ is such that any $\sigma_{d}$ incident to $\sigma_{k}$ has no vertices in $B_{\sigma_{k}}$, then $B_{\sigma_{k}}$ is contained in the union of the diametric balls of the Delaunay $d$-simplices incident to $\sigma_{k}$, and hence $\sigma_{k}$ is $k$-Gabriel.

Proof Since by definition, $B_{\sigma_{d}}$ is empty for any Delaunay simplex $\sigma_{d}$, it follows immediately that if $B_{\sigma_{k}}$ is contained in the union of diametric balls of Delaunay $d$-simplices incident to $\sigma_{k}$, then $\sigma_{k}$ is $k$-Gabriel.

Suppose to the contrary that $x \in B_{\sigma_{k}}$ and $x \notin B_{\sigma_{d}}$ for any $d$-simplex incident to $\sigma_{k}$. Point $x$ cannot lie in $\operatorname{aff}\left(S_{\sigma_{k}} \cap S_{\sigma_{d}}\right)$, since $B_{\sigma_{k}}$ and $B_{\sigma_{d}}$ coincide within that space. It follows that $x \notin \operatorname{aff}\left(\sigma_{k}\right)$. Therefore the segment $\left[x, c_{\sigma_{k}}\right]$ must intersect a face of some $d$-simplex, $\sigma_{d}$, incident to $\sigma_{k}$. Let the point of intersection be $y$.

Points $x$ and $y$ lie on the side of $\operatorname{aff}\left(S_{\sigma_{k}} \cap S_{\sigma_{d}}\right)$ in which $B_{\sigma_{k}}$ contains $B_{\sigma_{d}}$. Let $\sigma$ be the face of $\sigma_{d}$ that contains $y$. Then $\sigma$ has a vertex on the same side of $\operatorname{aff}\left(S_{\sigma_{k}} \cap S_{\sigma_{d}}\right)$ as $y$. But this is the side in which $B_{\sigma_{k}}$ contains $B_{\sigma_{d}}$, contradicting the hypothesis that $\sigma_{d}$ has no vertices in $B_{\sigma_{k}}$.

### 4.7 Discussion

The Gabriel property is a natural one to consider to define a manifold simplicial Delaunay structure. However, the observations of Section 4.5 demonstrate that the condition is too strong to be the keystone of a workable theory. In fact, the Gabriel perspective does not define a Delaunay structure in the sense used in Section 3.4: the Gabriel condition is not easily expressed as a Delaunay condition in terms of a choice of metric on $S$.

Almost Gabriel: Although Gabriel meshes suffer from theoretical existence problems, they have been exploited in practice. Presumably the violations of the Gabriel property are manifest by vertices interior to a diametric ball, but so close to the boundary that the problem has been mistaken for numerical error. We speculate that algorithms which have purported to construct Gabriel meshes have employed small tolerances designed to accommodate numerical error when checking the Gabriel condition. Thus what have been constructed in practice are "almost Gabriel" meshes. These structures have demonstrated their utility. It may be worthwhile to put them on a theoretical foundation.

Pseudo-disks: In Chapter 5 we will see that an important property for facilitating the construction of a Delaunay triangulation of $S$ is that the Delaunay circumdisks be pseudo-disks: their boundaries qualitatively intersect like Euclidean circles. In particular, the boundaries of two pseudodisk cannot intersect tangentially at more than one point, or intersect transversely at more than two points.

The pseudo-disk property of triangle circumcircles can also be seen to have relevance in the context of Delaunay edge flipping. Consider: if it is possible


Figure 4.7: Not pseudo-disks to define a metric that characterizes the Gabriel condition as a Delaunay condition, then the intersection of $B_{t}$ with $M$ would be the circumdisk of $t$ in that metric. Then triangle $t$ will yield a Gabriel certificate to $e$ if the vertex of the other triangle adjacent to $e$ does not lie within the circumdisk of $t$. Since this condition is not symmetric with respect to the two triangles adjacent to $e$, the triangle circumdisks do not have the pseudo-disk property, as evidenced by Figure 4.7.

We showed a specific obstruction to the construction of a Gabriel mesh that triangulates $S$, but perhaps the existence of asymmetric cases is a more general indication of a problem in a Delaunay-like criterion.

(a) Top view

(b) Side view

Figure 4.8: There can be no extrinsic edge flip algorithm to produce self-rDt meshes. We consider the flip-tet introduced in Section 4.5, drawn in blue here. In order to decide whether edge $e=[p, q]$ is locally self-rDt, we need to know whether triangle $[v, w, q]$ (drawn in red) passes above or below the circumcentre of the flip-tet. This is because the local Voronoi vertices for the hinge are given by the intersection with $M$ and the dual Voronoi edges in the 3D Voronoi diagram of the four vertices of the flip-tet in isolation. Here we have used $V(p v q)$ to denote the Voronoi edge dual to $[p, v, q]$.
self-rDt meshes: Chew [Che93] showed that the local Delaunay criterion for an rDt is symmetric about a shared edge when the triangles on $S$ are small enough. In that context he was considering edge flipping relative to the fixed geometry of $S$, but it is natural to wonder if the idea can be extended to yield an extrinsic edge flipping algorithm for producing self-rDt meshes.

In fact, no edge flipping algorithm which considers the local properties of a hinge alone can converge towards a self-rDt mesh. The problem is most easily illuminated from the Voronoi perspective. In general, a hinge $(u, e, v)$ on $e=[p, q]$ is not locally Delaunay if the circumdisk of $[u, p, q]$ does not contain $v$, and that of $[v, q, p]$ does not contain $u$. This can be equivalently formulated by considering the Voronoi diagram of the four vertices in isolation. Then $e$ is locally Delaunay only if the circumcentre of each of its adjacent triangles is a

Voronoi vertex in this simplified Voronoi diagram.
The problem for an extrinsic edge flip algorithm for self-rDt meshes is that the condition of whether the circumcentre of $t$ is a local Voronoi vertex cannot be decided on the basis of the properties of the hinge alone. Our obstruction to Gabriel meshes can be employed to demonstrate this. In the extrinsic metric, if the circumcentres of the triangles do not lie on the hinge, then information about the triangles outside the hinge is required to determine whether the "locally self-rDt" property is satisfied. The situation is illustrated in Figure 4.8.

Self-Delaunay mesh existence: The tetrahedral obstructions to the construction of Gabriel meshes are not an obstruction to self-Delaunay meshes: Lemma 3.10, tells us that at least one of the two hinges in a flip-tet must be locally Delaunay. However, the situation for self-Delaunay meshes is not free of difficulties. As we will discuss in Chapter 6, and then in more detail in Chapter 7, the Delaunay extrinsic edge flipping algorithm may encounter unflippable edges. An nlD edge $e$ is unflippable if its opposing edge $e^{\prime}$ already exists in the mesh. This issue implies that $P$ may not admit a self-Delaunay mesh homeomorphic to $S$. As a trivial example, the vertices of the tetrahedron in Figure 3.2, admit only one closed manifold triangle mesh, and it is not a self-Delaunay mesh.

The desire to find conditions on $P$ that will guarantee the existence of a self-Delaunay mesh that triangulates $S$ motivated the investigation into intrinsic sampling criteria that we present in Chapter 5. The closely related problem of unflippable edges plays a prominent role in the work presented in Chapters 6 and 7.

## Chapter 5

## Intrinsic surface sampling criteria

In this chapter we develop sampling conditions sufficient to ensure a well formed iVd, and hence that the iDt-mesh triangulates $S$. We utilize intrinsic sizing functions that were previously exploited by Leibon and Letscher [LL00], but we have relaxed their sampling criteria considerably. Our approach is inspired by the work of Boissonnat and Oudot [BO05], where Voronoi-Delaunay duality and the notion of pseudo-disks were both elegantly employed. Leibon [Lei99] also worked with the pseudo disk property, but the exposition is significantly more involved, and does not exploit the iVd.

We begin in Section 5.1 with an exposition of extrinsic and intrinsic sizing functions. Some of these functions were introduced briefly in Chapter 2. We make explicit the relationship between the various extrinsic sizing functions, and likewise for the intrinsic ones. There are thus two families of sizing functions, but the relationship between these families has not previously been quantified. In Section 5.3 we bridge this gap and provide explicit lower bounds on the intrinsic sizing functions with respect to the local feature size (the most important extrinsic sizing function). This allows us to compare our intrinsic sampling criteria, which we develop in Section 5.2, with previous sampling criteria for topological consistency which are based on the lfs.

### 5.1 Sizing functions for surface sampling

We present two families of sizing functions, each family being represented as a hierarchy of functions, with each individual function being bounded above or below by its neighbour in the hierarchy. The first family, which we call "extrinsic sizing functions" has four members,
all of which we've encountered in some form in Chapter 2. In preparation for introducing the "intrinsic sizing functions" in Section 5.1.3, we first discuss, in Section 5.1.2, the notion of strong convexity on a surface. In Section 5.3 .1 we will find inequalities that bridge the gap between representatives of these two families of sizing functions.

### 5.1.1 A natural hierarchy of extrinsic sizing functions

We recall here the definitions of quantities that were introduced in Chapter 2 with a specific focus on their mutual relationship as sizing functions.

Local feature size: The local feature size (lfs) at a point $x \in S$, denoted $\rho_{f}(x)$, is the distance from $x$ to the medial axis of $S$. It has become a de facto standard sizing function for extrinsic sampling criteria, and we will go to some lengths to compare our sampling criteria to those that are expressed in terms of the lfs.

The lfs enjoys the important property of Lipschitz continuity,

$$
\begin{equation*}
\left|\rho_{f}(x)-\rho_{f}(z)\right| \leq d_{\mathbb{R}^{3}}(x, z), \tag{5.1}
\end{equation*}
$$

allowing us to bound $\rho_{f}(x)$ in terms of a nearby $\rho_{f}(z)$.

Local reach: Each point $x \in S$ is associated with two medial balls, one on each side of the surface; one of them may have infinite radius. These balls are tangent to $S$ at $x$. The radius, $\rho_{R}(x)$, of the smaller of the two medial balls at $x$ is called the local reach at $x$. It is the distance from $x$ to the medial axis along a direction normal to $S$ at $x$. It was introduced by Federer [Fed59] where it was observed that it is a continuous function on


Figure 5.1: lfs vs. local reach. $S$. It is bounded below by the lfs: $\rho_{R}(x) \geq \rho_{f}(x)$ for all $x \in S$.

In Section 2.4 we exploited the local reach to define the tubular neighbourhood, $U_{\rho_{R}}$, of $S$ on which the orthogonal projection $\xi_{S}: U_{\rho_{R}} \rightarrow S$ is well defined: for $m \in U_{\rho_{R}}, \xi_{S}(m)$ is the closest point to $m$ on $S$. In the literature a tubular neighbourhood of $S$ is often defined in terms of the reach of $S: \rho_{R}=\inf _{x \in S} \rho_{R}(x)=\inf _{x \in S} \rho_{f}(x)$. However, for adaptive sampling
it is easier to guarantee that a mesh lies within the larger tubular neighbourhood that is defined by the local reach.

Maximal curvature: The maximal curvature at $x \in S$ is the maximum of the absolute values of the two principal curvatures: $\kappa(x)=\max \left\{\left|\kappa_{1}(x)\right|,\left|\kappa_{2}(x)\right|\right\}$. The associated sizing function, the maximal curvature radius, is given by the radius of the associated osculating sphere: $\rho_{\kappa}(x)=1 / \kappa(x)$.

Since the smallest medial ball at $x$ can never exceed the size of the osculating sphere at $x$, we have $\rho_{R}(x) \leq \rho_{\kappa}(x)$. Although the maximal curvature is continuous, it can vanish. Thus $\rho_{\kappa}(x)$ is not bounded nor even well defined everywhere. However, for its principle employment as an upper bound on the lfs, we may interpret $\rho_{\kappa}(x)=\infty$ at those points where the maximal curvature vanishes.

Gaussian curvature: The Gaussian curvature is the product of the principle curvatures: $G(x)=\kappa_{1}(x) \kappa_{2}(x)$. The Gaussian curvature radius, defined by $\rho_{G}(x)=1 / \sqrt{G(x)}$, is the associated sizing function. Like the maximal curvature radius, the Gaussian curvature radius suffers from a problem of definition when the Gaussian curvature is non-positive. Read $\rho_{G}(x)=\infty$ when $G(x) \leq 0$. We have $\rho_{\kappa}(x) \leq \rho_{G}(x)$.

The hierarchy: For all $x \in S$ we have the relations

$$
\begin{equation*}
\rho_{f}(x) \leq \rho_{R}(x) \leq \rho_{\kappa}(x) \leq \rho_{G}(x) . \tag{5.2}
\end{equation*}
$$

Of the four sizing functions represented here, the lfs is the one of primary interest. However, the other three functions play an important role in theoretical manipulations. The Gaussian curvature can actually be intrinsically defined, unlike the other three, but it fits comfortably into this hierarchy with the extrinsic functions.

### 5.1.2 Strong convexity

We seek a sampling criterion that guarantees a well formed Voronoi diagram and thus an iDt-mesh that triangulates $S$. To this end it is useful to examine the notion of convexity of sets on a surface. In the planar setting a set $A$ is convex if a line segment connecting any two points in $A$ lies in $A$. On a surface, lines are replaced by geodesics. There are several ways to extend the notion of convexity to sets on a surface. We follow Chavel [Cha06]:

Definition 5.1 (Strongly convex set on a surface) A set $A \subset S$ is strongly convex if for every $p, q \in A$,

1. there is a unique minimal geodesic $\gamma$ in $S$ connecting $p$ and $q$;
2. $\gamma$ lies entirely within $A$;
3. no other geodesic connecting $p$ and $q$ lies within $A$.

The set $A$ is convex if it satisfies only the first two conditions above, but we are interested in strong convexity. Note that there are many non-equivalent definitions of convexity and strong convexity in the Riemannian geometry literature, so care must be taken when referring to other works.

The intersection of strongly convex sets is strongly convex. Such sets are also contractible [dC92], which implies the following useful observation:

Lemma 5.2 A strongly convex set is simply connected.

### 5.1.3 Intrinsic sizing functions

We now present intrinsic sizing functions which embody standard concepts in Riemannian geometry. An introductory textbook on Riemannian geometry, e.g., [dC92], may be consulted for further details on the statements made here.

## Strong convexity radius

On the plane, convexity has to do with the shape of a set, but on a surface, the strong convexity condition also limits the size of the set. For example a geodesic disk is not strongly convex in general. Consider a geodesic disk on a cylinder. If the radius of the disk exceeds one quarter the circumference of the cylinder, then there will be points on the disk whose shortest connecting geodesic leaves the disk.

Definition 5.3 (Strong convexity radius) The strong convexity radius (scr) at a point $x \in S$ is defined as

$$
\rho_{\mathrm{sc}}(x)=\sup \left\{\rho \mid B_{S}(x ; r) \text { is strongly convex } \forall r<\rho\right\}
$$

It can be shown [dC92] that for any $x \in S, \rho_{\mathrm{sc}}(x)>0$. The scr is an intrinsic quantity, and in general, if $x$ is in a region of high curvature, $\rho_{\mathrm{sc}}(x)$ will be small. However, Gaussian curvature alone is not sufficient to characterize the scr. Consider again the example of a cylinder. The scr at any point will be no greater than $1 / 4$ the circumference of the cylinder, but the Gaussian curvature radius is unbounded.

There are no continuity results for the strong convexity radius. However, it is worth mentioning that Klingenberg [Kli95][1.9.9] has a definition of strong convexity in which a further axiom is imposed: A set $A$ is strongly convex in Klingenberg's sense if it satisfies the axioms of Definition 5.1 and further has the property that any geodesic disk $B$ contained in $A$ is also convex. It easily follows that the resulting scr is 1 -Lipschitz. It is shown that the scr is always positive even with this additional axiom.

## Injectivity radius

An explanation of the injectivity radius requires a brief description of the exponential map. We denote by $T_{x} S$ the tangent plane of $S$ at $x$. The exponential map at $x$ is a smooth mapping $\exp _{x}: T_{x} S \longrightarrow S$ that takes $X \in T_{x} S$ to the point $\gamma_{X}(\|X\|) \in S$, where $\gamma_{X}$ is the geodesic emanating from $x$ with tangent vector $X /\|X\|$. Restricted to a small enough disk in $T_{x} S, \exp _{x}$ is a diffeomorphism onto its image [dC92].

Definition 5.4 (Injectivity radius) The injectivity radius at $x \in S$ is the supremum of the radii for which $\exp _{x}$ is injective:

$$
\rho_{i}(x)=\sup \left\{\rho \mid \exp _{x} \text { is injective on } B_{T_{x} S}(0 ; r) \forall r<\rho\right\} .
$$

The function $\rho_{i}(x)$ is continuous on $S$ [Cha06]. One of the most useful properties of $\rho_{i}(x)$ follows from the definition: if $d_{S}(x, p)<\rho_{i}(x)$, then there is a unique minimal geodesic $\gamma$ between $x$ and $p$ and it will be the only geodesic between $x$ and $p$ that is contained in $B_{S}\left(x ; \rho_{i}(x)\right)$. By the third axiom of Definition 5.1, the radius of a strongly convex disk cannot exceed the injectivity radius of the centre. Thus $\rho_{\mathrm{sc}}(x) \leq \rho_{i}(x)$.

The image under $\exp _{x}$ of any disk $D$ centred at 0 in $T_{x} S$ is exactly a geodesic disk, and if the radius of $D$ is less than $\rho_{i}(x)$, then the image of $D$ will be a topological disk. However, it is incorrect to say that $\rho_{i}(x)$ is the largest radius for which the geodesic disk centred at $x$ is an embedded topological disk. Indeed, as the radius increases, there are two ways that
the exponential map can fail to be injective. One is if the disk wraps around and merges with itself to create nontrivial topology.

But the other situation that can occur is that the Jacobian of $\exp _{x}$ may become degenerate. Suppose this happens at a point $z$. The concentric geodesic circles centred at $x$ will be smooth provided their radius is less than $\rho_{i}(x)$, however the circle through $z$ may have a cusp. If $\gamma$ is the minimal geodesic joining $x$ to $z$ we say that $z$ is a


Figure 5.2: When $\exp _{x}$ is not injective conjugate point to $x$ along $\gamma$ (and vice versa: it is reciprocal). For all points $y$ on $\gamma$ between $x$ and $z$, the minimal geodesic between $x$ and $y$ will be a portion of $\gamma$. However, for all points on the extension of $\gamma$ past $z, \gamma$ will not be a portion of the minimizing geodesic [Cha06]. Thus $z \notin B_{S}\left(x ; \rho_{i}(x)\right)$. More generally, if an open neighbourhood $V$ of $x$ is the bijective image of $\left.\exp _{x}\right|_{U}$ for some $U \subset T_{x} S$, then $V$ contains no conjugate points and hence, by the inverse function theorem $\left.\exp _{x}\right|_{U}$ is a diffeomorphism onto $V$.

A geodesic loop is a geodesic that starts and ends at the same point. A closed geodesic (sometimes called a periodic geodesic) is a geodesic loop $\gamma$ whose tangent vectors agree at its endpoints: $\gamma^{\prime}(0)=\gamma^{\prime}(\ell(\gamma))$. Returning now to the case where a change in the topology of the image of $\exp _{x}$ occurs, it can be shown that there will be a geodesic loop starting and ending at $x$. The midpoint of this loop is the closest point to $x$ at which $\exp _{x}$ fails to be injective [Cha06].

If we extend a geodesic $\gamma$ from $x$, there will be a closest point $z \in \gamma$ beyond which $\gamma$ is no longer a minimizing geodesic. The set of all such points is called the cut locus of $x$ and it is compact [dC92]. The assertions made above are summarized by a theorem [Cha06][III.2.4, p.118] due to Klingenberg:

Lemma 5.5 (Klingenberg) If $q$ is the point on the cut locus of $x$ that is closest to $x$, then $q$ is either
(i) conjugate to $x$ along a minimal geodesic connecting them, or,
(ii) the midpoint of a geodesic loop starting and ending at $x$.

## Intrinsic sampling radius

There are no sampling criteria based exclusively on the injectivity radius. We give a sampling criterion, Corollary 5.9, that is based on the scr alone, but by combining the two sizing functions we are able to relax that criterion.

Definition 5.6 (Intrinsic sampling radius) The intrinsic sampling radius at $x \in S$ is given by

$$
\rho_{m}(x)=\min \left\{\rho_{\mathrm{sc}}(x), \frac{1}{2} \rho_{i}(x)\right\} .
$$

As with the extrinsic sizing functions, the family of intrinsic sizing functions is summarized by a hierarchy of inequalities:

$$
\rho_{m}(x) \leq \rho_{\mathrm{sc}}(x) \leq \rho_{i}(x) .
$$

### 5.2 Intrinsic sampling criteria

We now develop sampling criteria for a topologically consistent iDt-mesh. The sampling radii considered are based on the intrinsic sizing functions introduced in Section 5.1.3. Our analysis is in terms of the iVd, and we will generally implicitly assume the "intrinsic" modifier in reference to Voronoi structures in the remainder of this chapter.

### 5.2.1 The closed ball property via strong convexity

As the sampling density increases we expect the Voronoi cells to more closely exhibit the characteristics of those in a planar Voronoi diagram. One notable characteristic of a Voronoi cell in the plane is that it is convex. Voronoi cells on a surface cannot share this property, however, because Voronoi edges are not necessarily geodesics.

To see this, let $p, q$ lie on a plane, in which
 case their Voronoi boundary is a straight line (a geodesic). Now introduce a small bump near $p$, between $p$ and $q$, that is far from the original Voronoi boundary, as shown in the figure. This will distort the originally straight boundary so that it is no longer a geodesic. For a general surface, regardless of the sampling density, we
cannot expect Voronoi edges to be geodesics ${ }^{1}$, and therefore, we cannot demand convexity from the Voronoi cells on $S$ : two points on the boundary between two Voronoi cells cannot be connected by a minimal geodesic that lies within both cells. However, a sufficient sampling density will ensure that the Voronoi cells are contained in strongly convex neighbourhoods, which turns out to be a useful criterion.

Theorem 5.7 (Strong convexity and the iVd) If $\forall p \in P$ there exists a strongly convex set $U_{p} \subset S$ with $\mathcal{V}(p) \subset U_{p}$, then the iVd of $P \subset S$ is well formed.

The proof of Theorem 5.7 relies on the following lemma.

Lemma 5.8 Suppose that $\Omega \subset S$ is a union of Voronoi cells that is bounded by only two Voronoi cells, $\mathcal{V}(p), \mathcal{V}(q) \subset \overline{S \backslash \Omega}$, with $\mathcal{V}(p) \cap \mathcal{V}(q) \cap \Omega \neq \emptyset$. Then $\Omega$ contains a geodesic $\gamma$ that cannot be contained in a strongly convex set.


Figure 5.3: $\Omega$ is a region enclosed by the boundaries of $\mathcal{V}(p)$ and $\mathcal{V}(q)$. (a) The geodesic $\gamma$ is at least as long as the green curve and cannot be contained in any strongly convex set. (b) The minimal geodesic $\alpha$ must be contained in $U_{p}$. If $\alpha$ does not cross $\gamma$, then either $U_{q}$ or $U_{p}$ contains a loop of minimal geodesics (red and green curves) that encompasses $\gamma$. Assuming it is $U_{p}$, as shown, then $U_{p}$ cannot be simply connected and not contain $\gamma$. (c) If $\alpha$ does cross $\gamma$, then $U_{p}$ (or $U_{q}$ ) will contain two geodesics between $a$ and $z$.

Proof Let $a$ be a Voronoi vertex in $\mathcal{V}(p) \cap \mathcal{V}(q) \cap \Omega$, and let $s \in P$ be a sample in $\Omega$ with $a \in \mathcal{V}(s)$. Let $\gamma_{a s}$ be a minimal geodesic between $a$ and $s$. Then by Lemma 2.6, we have $\gamma_{a s} \subset \mathcal{V}(s)$.

[^7]We extend $\gamma_{a s}$ through $s$ until it exits $\Omega$ at $t$. Let $\gamma=\gamma_{a t}$ denote the resulting geodesic. Clearly, $\gamma \in \Omega$. Without loss of generality, let $t$ lie on the boundary of $\mathcal{V}(p)$; see Figure 5.3(a). Note that if $\gamma$ did not exit $\Omega$, it would exceed the diameter of $S$ and be too long to reside in a strongly convex set.

Both $a$ and $t$ lie on the boundary of $\mathcal{V}(p)$ and, by our choice of $s, d_{S}(a, p)=d_{S}(a, s)$, since $a$ is a Voronoi vertex. Also $d_{S}(t, p) \leq d_{S}(t, s)$. Thus $\gamma$ is at least as long as the path between $a$ and $t$ comprised of a minimal geodesic between $a$ and $p$ and a minimal geodesic between $p$ and $t$, shown as the green curve in Figure 5.3(a). It follows that $\gamma$ cannot be a unique minimal geodesic between $a$ and $t$ and so cannot be contained in any strongly convex set.

Proof of Theorem 5.7 Suppose that $\mathcal{V}(p)$ is not a topological disk. By Lemma 5.2, $U_{p}$ is simply connected. Thus a homotopically nontrivial loop in $\mathcal{V}(p)$ is homotopically trivial in $U_{p}$, and so $U_{p}$ must contain a region $\Omega$ that is exterior to but bounded on all sides by $\mathcal{V}(p)$. By an argument identical to the proof of Lemma 5.8 we see that such an $\Omega$ cannot be contained in a strongly convex set. Therefore $\mathcal{V}(p)$ must be a topological disk.

Now suppose that $\mathcal{V}(p)$ and $\mathcal{V}(q)$ meet at more than one distinct Voronoi edge. These cells then bound a region $\Omega$ as described in Lemma 5.8. Consider the strongly convex neighbourhoods $U_{p}$ and $U_{q}$. Since $\mathcal{V}(p) \subset U_{p}$ and $\mathcal{V}(q) \subset U_{q}$, the intersection $U_{p} \cap U_{q}$ must contain all Voronoi edges in $\mathcal{V}(p) \cap \mathcal{V}(q)$. Also, since $U_{p} \cap U_{q}$ is strongly convex the minimal geodesic, $\alpha$, between points $a$ and $b$ on distinct Voronoi edges (see Figure 5.3(b)) must lie in $U_{p} \cap U_{q}$.

Let $\gamma$ be the geodesic in $\Omega$ that was constructed in Lemma 5.8. Suppose that $\alpha$ does not cross $\gamma$. Then $\gamma$ must be contained in the region bounded by a loop of minimal geodesics (a geodesic triangle in fact) involving $\alpha$ and either $p$ or $q$. Assume it is $p$; refer to Figure 5.3(b). Now $U_{p}$ must contain $\alpha$, the minimal geodesic between $a$ and $p$, and the one between $p$ and $b$. However, by Lemma 5.8, $U_{p}$ cannot contain $\gamma$. We arrive at a contradiction to Lemma 5.2.

Thus $\alpha$ must cross $\gamma$. Let $z$ be the first such intersection that is encountered on a traversal of $\gamma$ starting at point $a$ and assume that this portion of $\gamma$ between $a$ and $z$ is contained in a region bounded by a loop of minimal geodesics involving $p$ (green curves) and $\alpha$ (red curve); see Figure 5.3(c). But now there are two geodesics between $a$ and $z ; U_{p}$ cannot contain this portion of $\gamma$, but it is forced to if it is to contain $\alpha$ and remain simply connected. Again a contradiction.

### 5.2.2 Sampling via the strong convexity radius

Theorem 5.7 ensures that if the Voronoi cells can be contained in strongly convex neighbourhoods, then an iDt-mesh that triangulates $S$ can be constructed. We can use the strong convexity radius to define a sampling radius that will guarantee the conditions of Theorem 5.7.

Corollary 5.9 If $\forall x \in S$ there exists a $p \in P$ such that $p \in B_{S}\left(x ; \frac{1}{2} \rho_{\mathrm{sc}}(x)\right)$, then the iVd of $P$ on $S$ is well formed.

Proof For any $\mathcal{V}(p)$, choose $z \in \mathcal{V}(p)$ that is at a maximal geodesic distance from $p$. For any $y \in \mathcal{V}(p)$, the triangle inequality yields

$$
\begin{aligned}
d_{S}(z, y) & \leq d_{S}(z, p)+d_{S}(p, y) \\
& \leq d_{S}(z, p)+d_{S}(z, p) \\
& \leq \rho_{\mathrm{sc}}(z) .
\end{aligned}
$$



Therefore $\mathcal{V}(p) \subset B_{S}\left(z ; \rho_{\mathrm{sc}}(z)\right)$ and Theorem 5.7 applies.
This sampling condition ensures more. If we choose $x$ to be a
Voronoi vertex, then the associated samples must all lie within $B_{S}\left(x ; \frac{1}{2} \rho_{\mathrm{sc}}(x)\right)$ and so there are unique minimal geodesics between them. Thus Corollary 5.9 can be strengthened to apply to intrinsic Delaunay triangulations:

Corollary 5.10 If $\forall x \in S$ there exists a $p \in P$ such that $p \in B_{S}\left(x ; \frac{1}{2} \rho_{\mathrm{sc}}(x)\right)$, then the iDt of $P$ on $S$ exists.

By contrast, Theorem 5.7 itself is not a priory sufficient to ensure that the iDt itself exists. Thus the above sampling condition is stronger than the condition imposed by Theorem 5.7. Compared to Corollary 5.10, the sampling criterion of [LL00] is more complicated; it cannot be expressed simply in terms of a single sampling radius. Also, it requires at least that there be a $p \in P$ such that $p \in B_{S}\left(x ; \frac{1}{5} \rho_{\mathrm{sc}}(x)\right)$.

### 5.2.3 A weaker criterion: the intrinsic sampling radius

One observation that comes up in the demonstration of Corollary 5.10 is that if we have a sampling criterion that demands only that a sample lies within the strong convexity radius
of any point on $S$, then we are guaranteed that there will be a unique minimal geodesic between any two samples that are Voronoi neighbours. So the question arises, in this case where the $\epsilon$ has been doubled from $\frac{1}{2}$ to 1 : Is the iVd still well formed? This would imply the existence of the iDt.

We have not obtained an affirmative answer to this question. However, we are able to ensure that the iVd is well formed provided the sampling radius is no larger than the intrinsic sampling radius, $\rho_{m}$. To facilitate this result, we borrow some terminology from Boissonnat and Oudot [BO05]:

Definition 5.11 (pseudo-disks) A family $\left\{B_{i}\right\}$ of topological disks on $S$ are pseudo-disks if for any two distinct disks $B_{i}$ and $B_{j}$, their boundaries either do not intersect, or they intersect tangentially at a single point, or they intersect transversely at exactly two points. Qualitatively pseudo-disks intersect each other in the manner expected of Euclidean disks:

Lemma 5.12 (Three circles) If $B_{1}, B_{2}$, and $B_{3}$ are pseudo-disks whose boundaries all intersect at $p$ and $q$, then one of the disks is contained in the union of the other two.

Proof It is sufficient to show that the boundary of one of the disks is contained in the union of the other two disks.

Let $C_{i}=\partial B_{i}$. Each of these circles is composed of two arcs joining $p$ and $q$. Choose a consistent orientation on the circles and consider the three arcs emanating from $p$. One of these arcs must be inside one of the two other disks, but outside of the other. Suppose this arc belongs to $C_{2}$ (blue in figure), and that it is inside $B_{1}$ and outside $B_{3}$. Since the intersections are transversal, the other arc on $C_{2}$ must be outside $B_{1}$, but inside
 $B_{3}$. Thus $C_{2}$ is contained in $B_{1} \cup B_{3}$.

The following lemma is an improvement and simplification of its namesake in [Lei99]:
Lemma 5.13 (Small circle intersection) For $x \in S$ and $r<\rho_{m}(x)$, the disks $B_{S}(x ; r)$ are pseudo-disks.

Proof Let $C_{x}=\partial B_{S}\left(x ; r_{x}\right)$ and $C_{y}=\partial B_{S}\left(y ; r_{y}\right)$ be geodesic circles with $r_{x}<\rho_{m}(x)$ and $r_{y}<\rho_{m}(y)$ and $x \neq y$. If $C_{x}$ and $C_{y}$ intersect tangentially at $z$, then by the Gauss
lemma [dC92][p.69], the minimal geodesics $\gamma_{x z}$ and $\gamma_{y z}$ connecting $x$ and $y$ with $z$ must have parallel tangent vectors at $z$. It follows that $x, y$ and $z$ all lie on a common geodesic $\gamma$.

If $\gamma_{x z}^{\prime}\left(r_{x}\right)=-\gamma_{y z}^{\prime}\left(r_{y}\right)$, then $\ell(\gamma)=r_{x}+r_{y}<\frac{1}{2}\left(\rho_{i}(x)+\rho_{i}(y)\right)$. It follows that either $x \in B_{S}\left(y ; \rho_{i}(y)\right)$, or $y \in B_{S}\left(x ; \rho_{i}(x)\right)$ and $\gamma$ is the unique minimal geodesic connecting $x$ and $y$. If there were another intersection at $w \neq z$, then $\alpha=\gamma_{x w} \cup \gamma_{y w}$ would be a path between $x$ and $y$ with $\ell(\alpha)=r_{x}+r_{y}$, contradicting the unique minimality of $\gamma$.

If on the other hand $\gamma_{x z}^{\prime}\left(r_{x}\right)=\gamma_{y z}^{\prime}\left(r_{y}\right)$, then either $x \in \gamma_{y z}$ or $y \in \gamma_{x z}$. Assume the former. Then $r_{y}=r_{x}+d_{S}(x, y)$. Now if there is another intersection at $w$, this same equality must apply and we conclude that $x \in \gamma_{y w}$. But a geodesic of length $r_{y}$ emanating from $y$ can only meet $C_{y}$ once, and assuming $d_{S}(x, y)>0$, there is only one such geodesic that contains $x$. Therefore we must have $w=z$. Thus if $C_{x}$ and $C_{y}$ intersect tangentially at $z$, there can be no other points of intersection.

Now suppose that $C_{x}$ and $C_{y}$ intersect transversely at $z$. Then $D=B_{S}\left(x ; r_{x}\right) \cap$ $B_{S}\left(y ; r_{y}\right) \neq \emptyset$. Because it is strongly convex, $D$ must be a single connected component bounded by an arc of $C_{x}$ and an arc of $C_{y}$. It follows that $C_{x}$ and $C_{y}$ must intersect transversely at another point distinct from $z$, and that there can be no further transverse intersections.

Lemmas 5.12 and 5.13 provide an obstruction to neighbouring Voronoi cells sharing more than two Voronoi vertices. This yields our main sampling result:

Theorem 5.14 The iDt of $P$ on $S$ exists if

$$
\forall x \in S \quad \exists p \in P \text { such that } p \in B_{S}\left(x ; \rho_{m}(x)\right) .
$$

Proof The sampling condition implies that there is a unique minimal geodesic between samples that are Voronoi neighbours. It remains to prove that the iVd is well formed.

Since $p$ lies within the strong convexity radius of each $x \in \mathcal{V}(p)$, there is an open neighbourhood $V$ of $\mathcal{V}(p)$ such that $V \subset B_{S}\left(p ; \rho_{i}(p)\right)$. It follows that $V$ is the diffeomorphic image under $\exp _{p}$ of some $U \subset T_{p} S$ (see Section 5.1.3). This implies that $\mathcal{V}(p)$ is contractible (use $\exp _{p} \circ t \cdot \exp _{p}^{-1}, t \in[0,1]$ ). Thus $\mathcal{V}(p)$ is a topological disk.

It remains to show that $\mathcal{V}(p)$ and $\mathcal{V}(q)$ cannot share more than a single Voronoi edge. If this were the case, we would have a region $\Omega$, as in Figure $5.3(\mathrm{a})$, that is bounded completely by $\mathcal{V}(p)$ and $\mathcal{V}(q)$. Suppose that the Voronoi vertices $a$ and $b$ were both on the boundary of $\mathcal{V}(s) \subset \Omega$. Consider the geodesic circles centred at $a$ and $b$ and with radius $d_{S}(a, s)<\rho_{m}(a)$
and $d_{S}(b, s)<\rho_{m}(b)$, respectively. Since $a$ and $b$ are Voronoi vertices, these two circles would have to intersect at $p$ and $q$ in addition to $s$, contradicting Lemma 5.13.

Suppose then that the Voronoi vertices $a$ and $b$ are on the boundaries of $\mathcal{V}(s) \subset \Omega$ and $\mathcal{V}(v) \subset \Omega$, respectively. By hypothesis, $\mathcal{V}(p)$ and $\mathcal{V}(q)$ share more than a single Voronoi edge. Therefore there are at least four Voronoi vertices in $\mathcal{V}(p) \cap \mathcal{V}(q)$. Let $c$ be such a Voronoi vertex, distinct from $a$ and $b$. Consider the three disks centred at these Voronoi vertices and with radii such that their respective three closest samples lie on the boundary. These three disks are pseudo-disks, and their boundaries intersect at $p$ and $q$. Thus by Lemma 5.12 one of these disks, say $B_{a}$, is contained in the union of the other two. However, $B_{a}$ has a third sample, $s$, on its boundary, contradicting the fact that all three disks must have empty interiors. Thus the iVd must be well formed.

Since the injectivity radius is never smaller than the scr, we are assured that the conditions imposed by Theorem 5.14 are at least as weak as those demanded by Corollary 5.10. For the comparison with the lfs which we develop next, the result of Theorem 5.14 is twice as good.

### 5.3 Relating extrinsic and intrinsic sampling criteria

In this section we examine the relationship between intrinsic sampling criteria based on the scr or the intrinsic sampling radius, and extrinsic ones based on the lfs. We show that for any $\epsilon>0$, there exists an $\epsilon_{f}>0$ such that any sample set $P$ that satisfies the extrinsic criterion:

$$
\begin{equation*}
\forall x \in S \quad \exists p \in P \text { such that } p \in B_{\mathbb{R}^{3}}\left(x ; \epsilon_{f} \rho_{f}(x)\right), \tag{5.3}
\end{equation*}
$$

will also satisfy the corresponding intrinsic criterion:

$$
\begin{equation*}
\forall x \in S \quad \exists p \in P \text { such that } p \in B_{S}\left(x ; \epsilon \rho_{m}(x)\right) . \tag{5.4}
\end{equation*}
$$

A first step in this direction is to form an estimate on the intrinsic sampling radius based on the lfs. In particular, we seek a constant $C$ such that

$$
\begin{equation*}
\rho_{m}(x) \geq C \rho_{f}(x) \text { for any } x \in S . \tag{5.5}
\end{equation*}
$$

Both $\rho_{f}(x)$ and $\rho_{m}(x)$ become smaller as the local maximal curvature becomes larger. However, $\rho_{f}(x)$ also becomes smaller when geodesically distant points of the surface become
close in the ambient space - a property that $\rho_{m}(x)$ does not possess. Thus although the lfs may be bounded above by the intrinsic sampling radius, we cannot hope to find a constant that would allow us to reverse the inequality in Equation (5.5). For this same reason, without further qualifications, we can never guarantee that any intrinsic sampling (5.4) will satisfy a given extrinsic criterion (5.3).

In Section 5.3.1, we obtain an estimate for the constant $C$ in Equation (5.5). In equation (5.3), distances to the sample set $P$ are measured in the ambient space, whereas the intrinsic conditions (5.4) are specified with respect to geodesic distances on the surface, which are larger in general. So the next step is to put an upper bound on the geodesic distance between a point $x \in S$ and a nearby point $p \in P \subset S$ in terms of the Euclidean distance between them. This is done in Section 5.3.2. Finally, in Section 5.3.3 we develop an explicit relationship between the $\epsilon$ of Equation (5.4) and the $\epsilon_{f}$ of Equation (5.3) and we use this relationship to compare the sampling criterion derived in Section 5.2.3 with more familiar extrinsic sampling criteria for meshing and surface reconstruction.

### 5.3.1 A lfs estimate on the intrinsic sampling radius

In this section we arrive at Theorem 5.19, one of the main results of this chapter. By producing a relationship (5.11) between the intrinsic sampling radius and the lfs, it opens the door for comparing intrinsic and extrinsic sampling criteria. We exploit curvature bounds established in the Riemannian geometry literature. Two facts, Lemmas 5.15 and 5.16, which respectively give insight into case (i) and (ii) of the Klingenberg Lemma 5.5, enable us to get an estimate on the injectivity radius. With the addition of a result from Chavel, Lemma 5.18, we obtain an estimate on the scr.

To tackle case (i) of Lemma 5.5, a theorem [Cha06][II.6.3, p.86] attributed to Morse and Schönberg states that if $q$ is conjugate to $p$ along a geodesic $\gamma$ and the Gaussian curvature along $\gamma$ is bounded above by $G$, then $\ell(\gamma) \geq \pi / \sqrt{G}$. In other words, $\ell(\gamma) \geq \pi \inf _{z \in \gamma} \rho_{G}(z)$, and since $\gamma$ is a compact set, there will be a point that attains the bound. Thus we can state the theorem in a more convenient form:

Lemma 5.15 (Morse, Schönberg) If a geodesic $\gamma$ connecting $p$ to $q$ contains a point conjugate to one of its endpoints, or if $p$ and $q$ are conjugate along $\gamma$, then

$$
\exists z \in \gamma \text { such that } \quad \rho_{G}(z) \leq \frac{\ell(\gamma)}{\pi}
$$

For case (ii) of Lemma 5.5 there is an extension to open curves of a famous theorem by Fenchel which we can exploit. Let $\gamma$ be a smooth space curve from $p$ to $q$. Let $v$ be the vector from $p$ to $q$ in $\mathbb{R}^{3}$. Denote by $\alpha$ and $\beta$ the angles $\angle\left(\gamma^{\prime}(0), v\right)$ and $\angle\left(\gamma^{\prime}(\ell(\gamma)), v\right)$ respectively, and let $k_{\gamma}(t)=\left\|\gamma^{\prime \prime}(t)\right\|$ be the curvature of $\gamma$ at $\gamma(t)$. Then the inequality, which is referred to in [Top06][p. 56] as the Fenchel-Reshetnyak inequality, states that

$$
\begin{equation*}
\int_{\gamma} k_{\gamma}(t) d t \geq \alpha+\beta \tag{5.6}
\end{equation*}
$$

The Fenchel-Reshetnyak inequality applies to non-closed curves, however, if we have $\gamma(0)=\gamma(\ell(\gamma))=p$, then we can break $\gamma$ into two pieces and obtain a curvature bound by applying Equation (5.6) to each piece.

Choose a point $q$ on $\gamma$ and let $\gamma_{1}$ be the portion of $\gamma$ from $p$ to $q$ and let $\gamma_{2}$ be the remaining portion from $q$ back to $p$. Denote the associated initial and final angles by $\alpha_{1}, \beta_{1}$ and $\alpha_{2}, \beta_{2}$ respectively. Then $\beta_{1}+\alpha_{2}=\pi$ and

$$
\begin{align*}
\int_{\gamma} k_{\gamma}(t) d t & =\int_{\gamma_{1}} k_{\gamma_{1}}(t) d t+\int_{\gamma_{2}} k_{\gamma_{2}}(t) d t \geq \alpha_{1}+\beta_{1}+\alpha_{2}+\beta_{2}  \tag{5.7}\\
& \geq\left(\alpha_{1}+\beta_{2}\right)+\pi
\end{align*}
$$

All we can assert about $\left(\alpha_{1}+\beta_{2}\right)$ is that it is not negative. Thus for an arbitrary loop curve $\gamma$ we have that the total curvature is at least $\pi$. It follows that there must be some point on $\gamma$ where the curvature is at least $\pi / \ell(\gamma)$. Applied to geodesics, this gives us the following:

Lemma 5.16 (Fenchel, Reshetnyak) On a geodesic loop $\gamma$,

$$
\exists z \in \gamma \text { such that } \quad \rho_{\kappa}(z) \leq \frac{\ell(\gamma)}{\pi} .
$$

We obtain an estimate on the injectivity radius:
Theorem 5.17 (Injectivity radius estimate) For all $x \in S$

$$
\rho_{i}(x) \geq\left(\frac{\pi}{2+\pi}\right) \rho_{f}(x)
$$

Proof In the first case of Lemma 5.5, we have from Lemma 5.15 and Equation (5.2) that there exists a point $z$ in $\bar{B}_{S}\left(x ; \rho_{i}(x)\right)$ with $\rho_{f}(z) \leq \rho_{G}(z) \leq \rho_{i}(x) / \pi$. Using the Lipschitz continuity of lfs (5.1), we have

$$
\begin{aligned}
\rho_{f}(x) & \leq \rho_{f}(z)+d_{S}(x, z) \\
& \leq \frac{\rho_{i}(x)}{\pi}+\rho_{i}(x),
\end{aligned}
$$

and so $\rho_{i}(x) \geq\left(\frac{\pi}{1+\pi}\right) \rho_{f}(x)$.
In the second case, we have by Lemma 5.16, a $z$ in $\bar{B}_{S}\left(x ; \rho_{i}(x)\right)$ with $\rho_{f}(z) \leq \rho_{\kappa}(z) \leq$ $2 \rho_{i}(x) / \pi$, and Lipschitz continuity yields $\rho_{i}(x) \geq\left(\frac{\pi}{2+\pi}\right) \rho_{f}(x)$. Taking the smaller of the two bounds gives our estimate on $\rho_{i}(x)$.

Our estimate on the scr is based on a theorem in [Cha06][IX.6.1, p.404] which gives a global lower bound for the scr in terms of global bounds on the injectivity radius and the Gaussian curvature. The following lemma is extracted from Chavel's proof:

Lemma 5.18 (Chavel) Let $x p q$ be a geodesic triangle consisting of geodesics $\gamma_{1}, \gamma_{2}$, and $\gamma_{3}$, connecting $q$ with $x, x$ with $p$, and $p$ with $q$ respectively. Suppose there are constants $\rho_{0}, r_{0}>0$ with $\rho_{0} \leq \rho_{i}(x)$, and such that $\rho_{G}(z) \geq r_{0}$ on $B_{S}\left(x ; \rho_{0}\right)$. If

$$
\sum_{i=1}^{3} \ell\left(\gamma_{i}\right)<\min \left\{2 \rho_{0}, 2 \pi r_{0}\right\}
$$

then

$$
\gamma_{3} \subset B_{S}(x ; \rho),
$$

where $\rho=\max \left\{d_{S}(x, p), d_{S}(x, q)\right\}$.
A discussion of Lemma 5.18 and its proof appears in Section C.2.
This result, together with Lemmas 5.15 and 5.16, yields the main result of this section:
Theorem 5.19 (scr estimate) For all $x \in S$,

$$
\rho_{\mathrm{sc}}(x) \geq\left(\frac{\pi}{4+3 \pi}\right) \rho_{f}(x)
$$

Proof Consider the geodesic disk of radius $r$ centred at $x$. There are three ways in which $B_{S}(x ; r)$ can fail to be strongly convex: There exist $p, q \in B_{S}(x ; r)$ such that either
(i) the minimizer $\gamma$ connecting $p$ and $q$ is not unique, or
(ii) in addition to $\gamma$, there is another geodesic $\alpha$ connecting $p$ and $q$ and contained in $B_{S}(x ; r)$, or
(iii) $\gamma$ is not contained in $B_{S}(x ; r)$.

Case (i) cannot happen if $d_{S}(p, q)<\max \left\{\rho_{i}(p), \rho_{i}(q)\right\}$, and possibility (ii) is eliminated if $B_{S}(x ; r) \subset B_{S}\left(p ; \rho_{i}(p)\right)$ for any $p \in B_{S}(x ; r)$. Thus we eliminate the first two cases if we ensure that $\rho_{i}(p) \geq 2 r$ for all $p \in B_{S}(x ; r)$.

If $\rho_{i}(p)<2 r$, then by the Klingenberg Lemma 5.5 either
(a) there is a $z \in B_{S}(p ; 2 r) \subset B_{S}(x ; 3 r)$ that is conjugate to $p$ along a minimizing geodesic, or
(b) there is a geodesic loop in $B_{S}(p ; 2 r)$.

In case (a), the Morse-Schönberg Lemma 5.15 gives us a $z \in B_{S}(x ; 3 r)$ with $\rho_{f}(z) \leq \frac{2 r}{\pi}$. The Lipschitz continuity of lfs yields $\rho_{f}(x) \leq \rho_{f}(z)+d_{S}(x, z) \leq \frac{2 r}{\pi}+3 r$. Thus

$$
\begin{equation*}
r \geq\left(\frac{\pi}{2+3 \pi}\right) \rho_{f}(x) \tag{5.8}
\end{equation*}
$$

In case (b), the Fenchel-Reshetnyak Lemma 5.16 yields a $z \in B_{S}(x ; 3 r)$ with $\rho_{f}(z) \leq \frac{4 r}{\pi}$. Again using Lipschitz continuity to bring the lfs bound to $x$, we obtain

$$
\begin{equation*}
r \geq\left(\frac{\pi}{4+3 \pi}\right) \rho_{f}(x) \tag{5.9}
\end{equation*}
$$

Thus if we ensure that $r$ is smaller than the bounds (5.8) and (5.9), then $\rho_{i}(p) \geq 2 r$ for all $p \in B_{S}(x ; r)$, and cases (i) and (ii) cannot happen. It remains to consider case (iii). For this we turn to Chavel's Lemma 5.18.

Consider the geodesic triangle $x p q$ consisting of minimal geodesics, with notation as in Lemma 5.18. By hypothesis we now have $\rho_{i}(x) \geq 2 r$, and also $\sum_{i=1}^{3} \ell\left(\gamma_{i}\right)<4 r$. Thus the conditions of Lemma 5.18 are satisfied with $\rho_{0}=2 r$, provided there is no $z \in B_{S}(x ; 2 r)$ with $2 \pi \rho_{G}(z)<4 r$. If there were such a $z$, it would imply $\rho_{f}(z)<\frac{2 r}{\pi}$, and the Lipschitz shuffle to $x$ yields

$$
\begin{equation*}
r>\frac{1}{2}\left(\frac{\pi}{1+\pi}\right) \rho_{f}(x) . \tag{5.10}
\end{equation*}
$$

By (5.9), the smaller of the three estimates, we have that $B_{S}(x ; r)$ is strongly convex whenever $r \leq\left(\frac{\pi}{4+3 \pi}\right) \rho_{f}(x)$, and we obtain the theorem by the definition of the scr.

Since the constant in the scr bound is less than half of that of the injectivity bound, we can use it also as a bound on the intrinsic sampling radius. Thus

Corollary 5.20 For all $x \in S$,

$$
\begin{equation*}
\rho_{m}(x) \geq\left(\frac{\pi}{4+3 \pi}\right) \rho_{f}(x) . \tag{5.11}
\end{equation*}
$$

### 5.3.2 Bounding geodesic lengths

Since Euclidean distances between two points never exceed the geodesic distances, $p \in$ $B_{S}(x ; \rho)$ implies $p \in B_{\mathbb{R}^{3}}(x ; \rho)$. However, we need to make claims about the containment of points within geodesic disks, given their presence within a Euclidean ball. Estimates on $d_{S}(x, p)$ relative to $d_{\mathbb{R}^{3}}(x, p)$, for $x, p \in S$ sufficiently close, are provided in the works of [MT04] and [HPW06]. We follow the terminology and notation of the former.

We exploit the projection mapping, $\xi_{S}: U_{\rho_{R}} \rightarrow S$ discussed in Section 5.1.1. The relative curvature, $\omega(m)$, at a point $m \in U_{\rho_{R}}$ is defined as

$$
\omega(m)=\frac{d_{\mathbb{R}^{3}}\left(m, \xi_{S}(m)\right)}{\rho_{\kappa}\left(\xi_{S}(m)\right)} .
$$

From the definition of $U_{\rho_{R}}, d_{\mathbb{R}^{3}}\left(m, \xi_{S}(m)\right) \leq \rho_{R}\left(\xi_{S}(m)\right)$, so Equation (5.2) gives $\omega(m) \leq 1$.
Suppose that $p \in B_{\mathbb{R}^{3}}\left(x ; \epsilon_{f} \rho_{f}(x)\right)$ and let $\left.I=\right] x, p[$ be the open Euclidean line segment between $x$ and $p$, and let $\omega=\sup _{m \in I} \omega(m)$. Then according to [MT04]:

$$
\begin{equation*}
\ell\left(\xi_{S}(I)\right) \leq \frac{1}{1-\omega} d_{\mathbb{R}^{3}}(x, p) . \tag{5.12}
\end{equation*}
$$

If $\gamma$ is a minimal geodesic between $x$ and $p$, then $d_{S}(x, p)=\ell(\gamma) \leq \ell\left(\xi_{S}(I)\right)$ gives us the needed bound.

Since $p \in S$, and $\xi_{S}$ takes $m$ to the closest point on $S$, we have, $\xi_{S}(m) \in B_{\mathbb{R}^{3}}\left(x ; \epsilon_{f} \rho_{f}(x)\right)$ for all $m \in I$, and in particular, $d_{\mathbb{R}^{3}}\left(m, \xi_{S}(m)\right) \leq \frac{1}{2} d_{\mathbb{R}^{3}}(x, p) \leq \frac{1}{2} \epsilon_{f} \rho_{f}(x)$. For the denominator of $\omega$ we have $\rho_{\kappa}\left(\xi_{S}(m)\right) \geq \rho_{f}\left(\xi_{S}(m)\right)$, and by the Lipschitz continuity of lfs (5.1), $\rho_{f}\left(\xi_{S}(m)\right) \geq\left(1-\epsilon_{f}\right) \rho_{f}(x)$. Thus $\omega \leq \frac{\epsilon_{f}}{2\left(1-\epsilon_{f}\right)}$. For the estimate to be usable, we need $\omega<1$, so we demand $\epsilon_{f}<2 / 3$.

Plugging this estimate into (5.12) together with $d_{\mathbb{R}^{3}}(x, p)<\epsilon_{f} \rho_{f}(x)$ yields the needed bound on the geodesic length:

Lemma 5.21 If $p \in S \cap B_{\mathbb{R}^{3}}\left(x ; \epsilon_{f} \rho_{f}(x)\right)$, with $\epsilon_{f}<2 / 3$, then $p \in B_{S}\left(x ; \tilde{\epsilon}_{f}(x)\right)$ for

$$
\tilde{\epsilon} \geq \frac{\epsilon_{f}\left(1-\epsilon_{f}\right)}{1-\frac{3}{2} \epsilon_{f}} .
$$

### 5.3.3 Extrinsic criteria meeting intrinsic demands

Equipped with Equation (5.11) and Lemma 5.21 we determine that if $P$ satisfies equation (5.3), then it will also satisfy equation (5.4) provided that

$$
\frac{\epsilon_{f}\left(1-\epsilon_{f}\right)}{1-\frac{3}{2} \epsilon_{f}} \leq \epsilon \frac{\pi}{4+3 \pi} .
$$

## Relating extrinsic and intrinsic sampling criteria <br> 

Figure 5.4: Sampling criteria from other works are compared. The $\epsilon$ for intrinsic sampling (5.4) is on the horizontal axis. The required $\epsilon_{f}$ for equation (5.3) is on the vertical axis.

Putting $C=\epsilon \frac{\pi}{4+3 \pi}$, we get

$$
\begin{equation*}
\epsilon_{f}^{2}-\left(1+\frac{3}{2} C\right) \epsilon_{f}+C \geq 0 \tag{5.13}
\end{equation*}
$$

an inequality that will be satisfied whenever $\epsilon_{f}$ is smaller than the smaller of the two positive roots.

For our sampling criteria of Theorem 5.14, $\epsilon=1$, yielding $C \approx 0.234$ and $\epsilon_{f} \leq 0.204$ is required. This compares well with existing lfs sampling requirements for topological consistency. For example, $\epsilon_{f} \leq 0.1$ is required in [AB98]. In [BO05], a loose $\epsilon$-sample is required to have $\epsilon \leq 0.091$. A loose $\epsilon$-sample only requires samples to lie within $B_{\mathbb{R}^{3}}\left(c ; \epsilon \rho_{f}(c)\right)$ when $c$ is a vertex of the rVd . According to Corollary 4.10 of that work, such a sampling will be an $\epsilon_{f}$-sampling in the sense of Equation (5.3) for $\epsilon_{f} \approx 0.161$. These comparisons are summarized in Figure 5.4.

### 5.4 Discussion

Through an analysis of the iVd, we improved upon the sampling criteria of [LL00, Lei99]. By deriving inequalities relating the injectivity radius and the strong convexity radius to the local feature size, we have enabled comparison between sampling criteria in the intrinsic and extrinsic domains. As described in Section C.3, we expect that the constant in Equation (5.11) can be improved from $\frac{\pi}{4+3 \pi}$ to $\frac{1}{2}\left(\frac{\pi}{1+\pi}\right)$. In any event, the indication is that sharper bounds may result from an intrinsic analysis even if an algorithm is based on an extrinsic model.

The recent geometric accuracy analysis of Dai et al. [DLYG06] is from the intrinsic viewpoint, using Leibon and Letscher's work [LL00] as the topological correctness foundation. The sampling conditions required by their main theorem (Theorem 3), involve a minimum amongst terms representing Leibon and Letscher's criterion and (larger) scaled extrinsic sizing functions. Thus it can be both relaxed and simplified in light of our work.

While we have shown that the iDt-mesh, $M$, triangulates $S$ if $P$ satisfies the intrinsic sampling radius, it may be that $M$ is not embedded. If we employ extrinsic sampling criteria which imply this intrinsic sampling criterion, then we expect that $M$ will be embedded because no triangle can intersect the medial axis. In fact, this embedding was taken for granted in [LL00, DLYG06] and was the reason for imposing lfs bounds on the sampling radius in addition to the scr bound. However, a formal verification that the iDt-mesh is embedded when lfs sampling criteria are met has not appeared. The homeomorphism proofs [AB98, ACDL00, Dey07], which apply to a substructure of $\mathcal{D}^{3}(P)$, do not apply a priori to the iDt-mesh, whose relationship to $\mathcal{D}^{3}(P)$ remains poorly understood.

The work of Leibon and Letscher [LL00] was summarized in the article published in conference proceedings, but detailed proofs of the theorems were left for a full article which was never published. The intrinsic sampling criteria described applied to higher dimensional manifolds, not just surfaces, but in the absence of complete demonstrations, it must be recognized that the higher dimensional problem is still unresolved.

There is therefore significant motivation to develop the results of this chapter for higher dimensions. However, although the statements of Theorems 5.7 and 5.14 make sense in higher dimensions, their proofs are not easily extended. On the other hand, extending the proofs of Theorems 5.17 and 5.19 may require no more than a reworking of Section 5.1.1.

## Chapter 6

## Constructing self-Delaunay meshes

In this chapter we present algorithms which take an arbitrary manifold triangle mesh as input, and convert it to a self-Delaunay mesh.

In Section 6.1 we present a geometry-preserving algorithm which produces a self-Delaunay mesh which is isometric to the input mesh. The algorithm works by vertex insertion on the mesh edges. Edges are split according to a method of concentric shells which was introduced by Ruppert [Rup95] in the context of Euclidean meshing. In that context the concentric shells method was a heuristic developed for dealing with sharp angles in the input structure. However, we show that in our context the method leads to a termination proof of the edge splitting algorithm. A disadvantage of this algorithm is that it can potentially add many vertices to the original mesh, and the distribution of the vertices reflects the initial triangulation more than it does the underlying geometry.

We then consider, in Section 6.2, the extrinsic edge flipping algorithm introduced in Section 3.2.1. We demonstrate that a Delaunay extrinsic edge flip on a non-planar hinge reduces the surface area of the mesh. Thus the algorithm is guaranteed to terminate. However, there is another issue: unflippable edges. An nlD edge $e$ is an unflippable edge if its opposing edge $e^{\prime}$ already appears in the mesh. If we were to flip such an edge, the mesh would no longer be manifold. Therefore, in order to guarantee that the edge flipping algorithm will terminate in a Delaunay edge, we may need to apply the geometry preserving algorithm if there remains unflippable nlD edges and no flippable ones.

The primary disadvantage of the extrinsic edge swapping algorithm is that it can produce unpleasing artifacts if the input mesh is coarsely sampled. The algorithm does not preserve the geometry, and this is most evident when a hinge with a large dihedral angle is flipped.

Thus to alleviate the weaknesses of both algorithms we propose a combination of the two. We do Delaunay extrinsic edge flips, but "feature edges" are not allowed to flip. We define a feature edge as any edge in the input mesh that has a dihedral angle greater than some threshold. If such an edge is nlD, it must be split, and the new edges that comprised portions of the original edge remain flagged as feature edges.

In addition, we introduce, in Section 6.3, a mesh decimation algorithm that takes a dense self-Delaunay mesh as input and produces coarser self-Delaunay meshes as output. Relying on constrained optimization [GW03] and a quadric-based error metric [GH97], the algorithm produces a series of self-Delaunay meshes, at multiple levels of detail, which approximate the input mesh. The resulting self-Delaunay meshes have a better triangle distribution than is obtained through the geometry preserving refinement. Thus the decimation algorithm can serve well as a post-processing step on the edge splitting algorithm.

These algorithms have been implemented and in Section 6.4 we examine mesh statistics and show images for several examples.

### 6.1 Geometry-preserving Delaunay remeshing



Figure 6.1: The original edge $[p, q]$ is split at $s$, creating two planar (blue) and two physical (red) edges incident to $s$.

### 6.1.1 Remeshing algorithm

The refinement proceeds by subjecting nlD edges to an edge split, which inserts a new vertex along an original edge of the input mesh and connects the newly inserted vertex with two vertices opposite to the current mesh edge being split, as shown in Figure 6.1. New vertices added during refinement are called split vertices to distinguish them from the original vertices of the mesh. An edge that has a non-zero dihedral angle (non-coplanar adjacent faces) and is an original mesh edge, or part of an original mesh edge is called a physical edge. If $e^{\prime}$ is a portion of a pre-existing edge $e$, we say that $e^{\prime}$ is embedded in $e$.

```
while mesh \(M\) contains an nlD edge \(e=[p, q]\) do
    if \(e\) is not a physical edge then
        Edge flip at \(e\).
    else
        if \(p\) is an original mesh vertex then
            Split \(e\) at \(S V(p, e)\).
        else
            Split \(e\) at \(S V(q, e)\).
        end if
    end if
end while
```

Algorithm 1: Geometry-preserving Delaunay remeshing.

Edges between coplanar faces are called planar edges. Only physical edges are split in our algorithm; planar edges may be flipped without affecting the geometry.

A naive edge splitting algorithm, using edge bisection for example, will not terminate in general, a problem example is illustrated in Figure 6.2. The kind of refinement problem we are facing has been studied in the planar setting, where one seeks a conforming Delaunay triangulation [ET93]. Inspired by Shewchuk's exposition of "The Quitter" meshing algorithm [She97][§3.7], our scheme, outlined in Algorithm 1, employs the method of concentric shells introduced by Ruppert [Rup95].


Figure 6.2: A naive Delaunay edge splitting algorithm will not terminate in general. Here we show on the left a fragment of an initial umbrella on vertex $v$, followed by snapshots of a progression of edge bisections. Blue edges are introduced when the central or outer radial edges are split, and red edges are introduced when the other two edges are split. The progression would continue indefinitely, always leaving nlD edges adjacent to $v$.

The idea is to split edges emenating from an original vertex $p$ only at distances that are multiples of a power of two from $p$. The split vertex, $S V(p, e)$, with $p$ a vertex of edge $e$, is defined to be the point $s$ along $e$ that is the closest to the midpoint $m$ of edge $e$, but such
that $|[p, s]|=2^{k} \delta$, for some (possibly negative) integer $k$. Formally,

$$
\begin{equation*}
S V(p, e)=\operatorname{argmin}_{s \in e,|[p, s]|=2^{k} \delta, k \in \mathbb{Z}}\|s-m\| . \tag{6.1}
\end{equation*}
$$

The constant factor $\delta \in \mathbb{R}^{+}$may be any positive number. In our implementation, we set it to unity.


Figure 6.3: Geometry-preserving Delaunay remeshing on a simple example. Physical edges are shown in red and planar edges in blue.

Our Delaunay remeshing scheme combines edge flipping with edge splitting. Since only edges interior to the planar faces of the mesh can be flipped, the algorithm is geometrypreserving. Figure 6.3 shows our remeshing algorithm at work on a simple example.

### 6.1.2 Termination proof

It suffices to prove that only a finite number of split vertices will be added, since all our edge flips are planar and they terminate as in the planar case [Law77]. We assume that there are no degenerate triangles in the mesh. The proof is by contradiction. Since a mesh surface is compact, the Bolzano-Weierstrass theorem implies that any non-terminating refinement must produce an accumulation point, $a$. That is, any arbitrarily small neighbourhood of $a$ contains infinitely many inserted split vertices. The algorithm never inserts vertices interior to the original mesh faces. Thus there are only two possible cases: either $a$ lies in the interior of an original edge, $e$, or $a$ is an original vertex. We eliminate these two possibilities with Lemmas 6.2 and 6.3 below.

We first observe that the concentric shells split scheme guarantees that edges are split into pieces of size bounded by a fraction of the original edge.

Lemma 6.1 Suppose $e=[p, q]$ is split at $s=S V(p, e)$, creating a child edges $e^{\prime}=[p, s]$
and $e^{\prime \prime}=[s, q]$. Then

$$
\left|e^{\prime}\right|<\frac{2}{3}|e| \text { and }\left|e^{\prime \prime}\right|<\frac{2}{3}|e| .
$$

Proof The result is a consequence of the fact that the split vertex defined by Equation (6.1) is within a distance of $\frac{1}{6}|e|$ from the midpoint of $e$.

Suppose to the contrary that $\left|e^{\prime}\right|>\frac{2}{3}|e|$. Let $m$ be the midpoint of $[p, q]$ and $r$ the midpoint of $[p, s]$, as shown in Figure 6.4(a). Then

$$
|[r, m]|=\frac{|e|}{2}-\frac{\left|e^{\prime}\right|}{2}<\frac{|e|}{6} .
$$

But since

$$
|[s, m]|=\left|e^{\prime}\right|-\frac{|e|}{2}>\frac{2|e|}{3}-\frac{|e|}{2}=\frac{|e|}{6},
$$

we have $|[r, m]|<|[s, m]|$. Thus, according to equation (6.1), point $r$, also a "power of 2 split" since $|[p, r]|=\frac{1}{2}\left|e^{\prime}\right|$, should have been chosen as the split vertex instead of $s$.

Similarly, if $\left|e^{\prime \prime}\right|>\frac{2}{3}|e|$, then we would have $\left|e^{\prime}\right|<\frac{1}{3}|e|$. Choose point $r^{\prime}$ such that $\left|\left[p, r^{\prime}\right]\right|=2\left|e^{\prime}\right|$. Again $r^{\prime}$ is a valid candidate point for a split, and it is closer to $m$ than is $s$, contradicting the claim that $s$ is a split vertex.

Lemma 6.2 Split vertices do not have an accumulation point in the interior of an original mesh edge.

Proof Suppose $a$ is an accumulation point interior to an original edge $e$. Then there is an $\epsilon>0$ such that the geodesic disk $O$ of radius $\epsilon$ centred at $a$ contains none of the original vertices nor any portion of an original edge other than $e$. As edge splits happen along $e$, consider the sequence $e_{0}, e_{1}, \ldots$ of mesh edges which contain $a$, where $e_{0}=e$. Since $a$ is an accumulation point, this edge sequence must be infinite.

According to Lemma 6.1, $\left|e_{i}\right| \leq \frac{2}{3}\left|e_{i-1}\right|$ for all $i>0$. Thus we conclude that there must be some $e_{i}$ that falls entirely inside the geodesic disk $O$. Clearly, such an edge $e_{i}$ cannot possibly be nlD since barring degeneracy, its two opposite vertices must lie outside the disk $O$, as shown in Figure 6.4(b). Therefore, $e_{i}$, which contains $a$, will not be split further. Thus $a$ cannot be an accumulation point.

Lemma 6.3 Split vertices do not accumulate at an original vertex.


Figure 6.4: Figures for termination proof. (a) If $\left|e^{\prime}\right|>\frac{2}{3}|e|$, then $r$, the midpoint of $e^{\prime}$, should have been the split vertex, instead of $s$. (b) A geodesic disk $O$ centred at $a$ contains an edge covering $a$. Since the vertices opposite to the edge lie outside $O$, the edge cannot be nlD. (c) Eventually, all edge splits towards $a$ must occur on concentric geodesic circles, centred at $a$, with radii $2^{k} \delta$, where $k \in \mathbb{Z}$. (d) If $|[a, d]| \geq|[a, c]|$ and $|[a, g]| \geq|[a, c]|,[a, c]$ cannot be nlD.

Proof Suppose the original vertex $a$ is an accumulation point. Then there is an infinite accumulation of edge splits towards $a$. Eventually, such splits must all occur on geodesic circles centred at $a$ with radii $2^{k} \delta, k \in \mathbb{Z}$, as shown in Figure 6.4(c). Consider such a geodesic circle, $O_{1}$, that is sufficiently small that all original mesh edges meeting at $a$ extend beyond the next two concentric geodesic circles, $O_{2}$ and $O_{3}$. Suppose that $b$, along some edge $e$, is the very first split vertex that is created on $O_{1}$, as shown in Figure 6.4(d). It follows that $b$ must bisect some nlD edge $[a, c]$ embedded in $e$, where $c$ lies on $O_{2}$.

However, we can argue that edge $[a, c]$ cannot possibly be nlD, since the lengths of edges $[a, d]$ and $[a, g]$ in the quad containing vertices $a$ and $c$ must be greater than or equal to $|[a, c]|$; this is due to our assumption that $b$ is the first split vertex created inside $O_{2}$. To see that $|[a, d]|,|[a, g]| \geq|[a, c]|$ implies that $[a, c]$ is not nlD, note that since $[a, c]$ is not the longest edge in either of its adjacent triangles, it cannot subtend an obtuse angle.

Lemmas 6.2 and 6.3 together show that spit vertices cannot accumulate on the mesh. It follows that the refinement algorithm must terminate.

Theorem 6.4 The geometry preserving Delaunay refinement algorithm, Algorithm 1, terminates in a self-Delaunay mesh.

### 6.1.3 Meshes with boundaries

In [BS07, FSBS06] meshes with boundaries are dealt with by using a constrained Delaunay triangulation; there was not really another choice available in that context. However, with refinement, we are by necessity allowed to add Steiner vertices. Therefore, we can also ensure that all the boundary edges are Delaunay, where we define a boundary edge $e$ to be Delaunay if and only if it subtends a nonobtuse angle. This ensures that the edge would connect Voronoi neighbours and that it is contained in an empty geodesic disk.

Edge splitting at mesh boundaries works in exactly the same way as before. To show that an accumulation point cannot occur on the interior of an edge, note that if an edge embedded in the mesh boundary falls entirely inside the geodesic disk $O$, then it cannot be nlD because its opposite vertex lies outside $O$ and thus must form a nonobtuse angle (see Figure 6.4(b)). In the proof of Lemma 6.3, if edge $[a, c]$ happens to lie on the mesh boundary, then $|[a, d]| \geq|[a, c]|$ implies that $\angle a d c \leq \angle a c d$ and thus $\angle a d c \leq \pi / 2$. It follows that edge $[a, c]$ is locally Delaunay.

### 6.2 Delaunay remeshing via mesh edge flipping

The geometry-preserving refinement algorithm preserves geometry at the expense of adding vertices to the original mesh. In that algorithm, there are no flips of physical edges, but a self-Delaunay mesh can also be obtained by allowing physical edges to be flipped, using the extrinsic edge flipping algorithm as described in Section 3.2.1. Such operations do not introduce vertices, although they do compromise the geometry. However, due to the possibility of unflippable edges, a refinement step may be necessary to ensure that the final mesh is self-Delaunay. Flipping the physical edges of a mesh changes the nature of the problem considerably, for now the geometry of the underlying domain is changing each time an edge is flipped. The arguments for termination of edge flipping in the planar case [Mus97] no longer apply, and we also need to deal with possible topological constraints manifest in unflippable edges.

### 6.2.1 Edge flipping and refinement algorithm

Our algorithm takes any manifold triangle mesh $M$ as input and it is superficially similar to the one described by Bobenko and Springborn [BS07]. However, we may require an


Figure 6.5: A Delaunay extrinsic edge flip may decrease minimum angle and increase harmonic index: Let $e$ be nlD, then flipping $e$ into $e^{\prime}$ is a Delaunay flip. However, one can bend $\triangle[u, q, v]$ towards $\triangle[u, p, v]$ so that $\angle p u q$ becomes arbitrarily small and the Harmonic index of $\triangle[p, u, q]$ becomes arbitrarily large.
additional refinement step: a mesh is always associated with a proper triangulation. In this section we address the issues of termination and correctness of the extrinsic edge flipping algorithm, Algorithm 2.

In our implementation, the order in which the edges are flipped is realized using a priority queue, with the cost function being the sum of the angles subtended by the edge, minus $\pi$. However, there is currently no theoretical or experimental motivation for choosing this particular order over any other.
while mesh $M$ contains a flippable nlD edge $e$ do
Edge flip at $e$.
end while
Run geometry-preserving Delaunay remeshing.
Algorithm 2: Geometry-altering Delaunay remeshing

### 6.2.2 Delaunay extrinsic edge flipping and area minimization

Termination of edge flipping is traditionally shown by defining a functional on a triangulation and proving that it is increased (or decreased) with each edge flip. However, most of the traditional measures applicable to the case of fixed geometry do not extend to the case of Delaunay extrinsic edge flips. For example, the minimal angle in the triangles adjacent to an edge can sometimes be decreased after a Delaunay flip. Likewise, the harmonic index (sum of the cotangents of the angles), exploited by [BS07], may increase. See Figure 6.5 for an example. It turns out that a measure that is consistently non-increasing with each

Delaunay flip is the mesh surface area:

Theorem 6.5 If the sum of the two angles opposite to an edge $e$ is greater than the corresponding angle sum for the opposing edge $e^{\prime}$, then the area of the hinge associated with $e$ is greater than or equal to the area of the opposing hinge associated with $e^{\prime}$. Equality arises only when $e$ and $e^{\prime}$ lie in the same plane.

Proof We are concerned with the area of the quadrilaterals defined by the hinges. Each hinge can be made planar without distorting the area by unfolding it on the associated edge. The edges $e$ and $e^{\prime}$ define two different quadrilaterals, but they share the same set of sides. Let $a, b, c$, and $d$ be the lengths of each of the sides. We exploit Bretschneider's formula [Bre42] for the area of a quadrilateral:

$$
\begin{equation*}
\mathcal{A}=\sqrt{(s-a)(s-b)(s-c)(s-d)-a b c d \cos ^{2}\left(\frac{A+C}{2}\right)} \tag{6.2}
\end{equation*}
$$

where $s=(a+b+c+d) / 2$ is the semi-perimeter and $A$ and $C$ are angles opposite edge $e$. Let $B^{\prime}$ and $D^{\prime}$ be the angles opposite edge $e^{\prime}$ in the other quadrilateral.

Noting that $\cos ^{2} \theta$ is monotonically decreasing in the interval $[0, \pi / 2]$ and that $\frac{A+C}{2}>$ $\frac{B^{\prime}+D^{\prime}}{2}$ by hypothesis, we have

$$
\cos ^{2}\left(\frac{A+C}{2}\right)<\cos ^{2}\left(\frac{B^{\prime}+D^{\prime}}{2}\right), \text { if } \frac{A+C}{2}<\frac{\pi}{2} .
$$

Thus, by the Bretschneider's formula (6.2), the area of the quadrilateral associated with $e^{\prime}$ is less than the area of that associated with $e$. On the other hand, if $\frac{A+C}{2} \geq \frac{\pi}{2}$, then by equation (3.1), we have

$$
\pi / 2 \geq \pi-\frac{A+C}{2} \geq \frac{B^{\prime}+D^{\prime}}{2}
$$

with equality in the planar case. Thus

$$
\cos ^{2}\left(\frac{A+C}{2}\right)=\cos ^{2}\left(\pi-\frac{A+C}{2}\right) \leq \cos ^{2}\left(\frac{B^{\prime}+D^{\prime}}{2}\right)
$$

Again, the Bretschneider's formula gives us a decrease in area except when $e$ and $e^{\prime}$ lie in the same plane, in which case the area is unchanged.

Thus, the surface area of the mesh is monotonically non-increasing as we perform edge flips. Since the number of possible triangulations is finite, these edge flips either terminate, or become stuck in an endless sequence of planar edge flips. This latter possibility is eliminated


Figure 6.6: Edge $e$ is unflippable, because its opposing edge, $e^{\prime}$, already exists in the mesh. The flip-tet is outlined in red. (a) 2-exposed. (b) 3-exposed.
by traditional termination proofs of Delaunay edge flipping [Law77, BS07], i.e., the harmonic index of the mesh decreases with each flip in this case, for example. Thus the edge flipping step of Algorithm 2 terminates.

Next, we discuss the possibility of having nlD edges which still remain after termination of edge flipping. These edges could not be flipped due to topological constraints.

### 6.2.3 Unflippable edges

An edge $e$ in mesh $M$ is unflippable if its opposing edge also exists in $M$. Recall that the opposing edge to $e$ is the edge $e^{\prime}$ that would replace $e$ after a flip. If $e^{\prime}$ is already in the mesh, flipping $e$ would result in a non-manifold edge.

Examining the flip-tet $\sigma$ associated with $e$, we identify only three possible cases of unflippable edges. At least two of the faces of $\sigma$, those adjacent to $e$, belong to $M$. We say $\sigma$ is 2-exposed, 3-exposed, or 4 -exposed, corresponding to it having zero, one, or both of the remaining two faces belonging to $M$. The 4-exposed case is trivial as it can only occur when $M$ is a tetrahedron (assuming $M$ is connected). Figure 6.6 depicts the other two cases: 2 -exposed and 3 -exposed. The 3 -exposed case occurs when an nlD edge has a valence three vertex at one of its ends.

In Section 7.1 we show that unflippable nlD edges do not exist in a smooth mesh. However, this alone is not sufficient to characterize the meshes which can be flipped to a
self-Delaunay mesh without the addition of extra vertices; Delaunay extrinsic edge flipping may make an initially smooth mesh become non-smooth. These issues are explored in Chapter 7.

The problems presented by unflippable edges, as well as boundary edges, can be treated by the final geometry-preserving remeshing step in Algorithm 2, which ensures that all the edges in the final mesh are locally Delaunay.

Remark 6.6 In the original publication of the results of this chapter [DZM07a], we gave a heuristic argument to support the claim that unflippable edges were a symptom of a poorly sampled point set $P$. We argued that if the only nlD edges in a mesh are unflippable, then the iVd of the vertices of the mesh is not well formed. However, there were errors in the argument. In particular, we claimed that if the hinge $(u, e, v)$ is nlD, then $\mathcal{V}(u)$ and $\mathcal{V}(v)$ must meet on that hinge. We also made the claim that the opposing edge to an unflippable edge must be flippable. Both claims are incorrect.

We do not know whether or not a mesh whose only nlD edges are unflippable must have an iVd that is not well formed. However, even if this were the case, it does not immediately follow that a sufficiently dense sampling of $S$ is all that is needed to avoid the problem. As demonstrated in Chapter 5, if $P$ is sufficiently dense, the iVd of $P$ on $S$ will be well formed. It can be expected that if $P$ are the vertices of $M$, and $M$ is a good approximation of $S$, then the iVd of $P$ on $M$ will also be well formed. However, without guarantees that Delaunay edge flipping maintains some kind of error bound between $d_{M}$ and $d_{S}$, we cannot be assured that the closed ball property will be preserved during edge flipping.

It may well be that even if $P$ is well sampled from $S$, unflippable edges may be generated during Delaunay extrinsic edge flipping if the initial mesh is "too far" from being a selfDelaunay mesh. This problem of unflippable edges has motivated much of the research presented in this thesis.

### 6.3 Delaunay mesh decimation

For some applications it may be desirable to have a progression of self-Delaunay meshes, at multiple levels of details (LODs), which approximate an initial dense self-Delaunay mesh. In practice any resource-constrained application can benefit from LOD representations to avoid redundancy and allow for more effective use of the triangle budget. In this section
we describe a mesh decimation algorithm for LOD modelling with self-Delaunay meshes. Our geometry-preserving refinement, Algorithm 1, may produce an excess of small triangles, but the resulting self-Delaunay mesh serves well as input for the decimation algorithm we describe here. The output of this algorithm has a much more pleasing triangle distribution, but of course it is not geometry preserving.


Figure 6.7: An edge collapse. The edges affected (blue and red) need to remain locally Delaunay.

The Delaunay mesh decimation algorithm has been adapted from a similar scheme for nonobtuse meshes [LZ06] and is based on edge collapse prioritized by a quadric error as introduced by Garland and Heckbert [GH97]. For each edge collapse, the resulting vertex needs to lie in a feasible region to ensure that all the affected edges remain locally Delaunay. The optimal position of the vertex is chosen to minimize the standard quadric error [GH97], subject to constraints, and the resulting error sets the priority. By choosing the allowable region to be a linearized and convexified subset of the feasible region, i.e., being conservative, the decimation algorithm is formulated as a constrained least-square problem and is solved using the OOQP solver of Gertz and Wright [GW03]. Thus the error quadric associated with each edge collapse is minimized subject to linear constraints in the form of planes bounding half spaces.

Since quadric-based edge collapsing schemes using priority queues are well-known [GH97, LZ06], we will concentrate on describing the feasible region and its linearization. In Figure 6.7, we show an edge collapse $[u, v] \rightarrow w$ after which we need to ensure that all the incident edges at $w$ (shown in red), e.g., $[w, p]$ and $[w, q]$, and all the subtending edges to $w$ (shown in blue), e.g., $[p, q]$, remain locally Delaunay.

Allowable region associated with a subtending edge: Let $[p, q]$ be a subtending edge, and let $w^{\prime}$ be the other subtended vertex. Then any $w$ with $\angle p w q \leq \pi-\angle p w^{\prime} q$ is feasible. Consider the circumcircle of $\left[p, w^{\prime}, q\right]$. If we rotate the arc subtended by edge $[p, q]$ about $[p, q]$, as shown in Figure 6.8(a), we obtain a surface of revolution which we call a chordal spheroid. Any $w$ outside of the spheroid is feasible. To linearize the feasible region,


Figure 6.8: Feasible and allowable regions for subtending edges. (a) If $e=[p, q]$ is a subtending edge, the feasible region is the complement of the chordal spheroid obtained by rotating the subtended arc of the circumcircle of $\left[p, w^{\prime}, q\right]$, where $w^{\prime}$ is the other subtending vertex of $e$. (b) The allowable region is obtained by linearizing the feasible region by means of a tangent plane at $r$, as described in the text.
we replace it by a half space defined by a plane $L$ tangent to the sphere or the chordal spheroid defined above (Figure 6.8(b)). The point of tangency, $r$, is chosen so that $L$ is parallel to $[p, q]$ and perpendicular to aff $\left(\left[p, w^{\prime}, q\right]\right)$.

Allowable region associated with an incident edge: To ensure that an incident edge $[w, p]$ remains locally Delaunay, we need $\angle p o w+\angle p q w \leq \pi$, where $o$ and $q$ are the one-ring vertices of $w$ adjacent to $p$; see Figure 6.9(a).

An example of the actual surface bounding the feasible region in 3D is shown in Figure 6.9(b). To construct the linearized allowable region, we focus our attention on the circumcircle of $[o, p, q]$. If $w$ were confined to lie in the plane supporting $[o, p, q]$, then when $w$ lies within the wedge defined by $\angle o p q,[p, w]$ would be nlD if $w$ is outside the circumcircle of $[o, p, q]$. Therefore at a minimum our constraint planes must bound the allowable region away from this portion of the wedge.

We will call aff $([o, p, q])$ the horizontal plane. Our constraint planes will all be perpendicular to this plane, i.e., vertical. The constraint planes define half-spaces, the intersection of which is the allowable region. We focus our attention on the planes themselves: The allowable half-space will always be the one that contains the point $p$.


Figure 6.9: Feasible and allowable regions for incident edges. (a) For $[w, p]$ to be locally Delaunay, we need $\angle p o w+\angle p q w \leq \pi$. (b) An example plot of the actual surface bounding the feasible region with respect to $o, p$, and $q$. The two cusps on the surface correspond to $o$ and $q$, and the supporting plane of $[o, p, q]$ is parallel to the top face of the box. (c) Planes $L_{1}$ and $L_{2}$ enclose a conservative, linearized allowable region for our optimization in the simplest case. The definition of $L_{1}$ and $L_{2}$ is given in the text.

Let $p^{*}$ be the point antipodal to $p$ in the circumcircle of $[o, p, q]$. An allowable region may be simply described by the two vertical planes that contain the segments $\left[0, p^{*}\right]$, and $\left[q, p^{*}\right]$ respectively. Since these planes are orthogonal to $[p, o]$ and $[p, q]$ respectively, both $\angle p o w$ and $\angle p q w$ are acute in this region, and therefore it is necessarily contained within the feasible region. However, our experiments revealed these constraints to be too restrictive. On some models we would obtain clusters of vertices where no decimation had occurred.

Instead, we choose two initial planes which maximize the area of the allowable region within the circumcircle. We select $r$ to be the midpoint on the arc subtending $\angle o p q$. This selection maximizes the area of $[o, q, r]$. We now consider the two vertical planes $L_{1}$, and $L_{2}$, containing $[r, o]$ and $[r, q]$ respectively, as shown in Figure 6.9(c).

As described in Section D.2, $L_{1}$, and $L_{2}$ are not always sufficient to bound the allowable region within the feasible region. In order to construct correct allowable regions we distinguish between the opposite-side case, where $o$ and $q$ lie on opposite sides of the diameter through $p$, and the same-side case, where $o$ and $q$ both lie on the same side of $\left[p, p^{*}\right]$. If $o$ or $q$ happens to coincide with $p^{*}$, we treat it as an opposite-side case.


Figure 6.10: Allowable regions for incident edges. The shaded region is the allowable region for $w$ in the plane of $[o, p, q]$. The constraint planes are defined by vertical (i.e. perpendicular to aff $([o, p, q]))$ planes containing the red lines. The three cases are described in the text.

The circumcylinder is the vertical right cylinder defined by the circumcircle of $[o, p, q]$. If $[o, p, q]$ defines an opposite-side case, then within the circumcylinder, planes $L_{1}$ and $L_{2}$ contain the allowable region within the feasible region. However, these two planes may not entirely lie within the feasible region outside of the circumcylinder. Sometimes a third plane must be added.

The close vertex is whichever of $o$ and $q$ is closest to $p^{*}$, the other is the far vertex. We label them $c$ and $f$ respectively. If $\left|\left[c, p^{*}\right]\right|<\left|\left[p^{*}, r\right]\right|$, then we add a third vertical plane, $L_{t}$, which contains $c$ and $a$, where $a$ is the midpoint of $\widehat{c r}$, the arc between $c$ and $r$. Thus constructed, $L_{t}$ is sufficient to contain the allowable region within the feasible region. As discussed in Appendix D, $L_{t}$ may be viewed as a tangent plane at $c$ to the portion of the boundary of the feasible region that lies outside of the circumcylinder. (The boundary of the feasible region has a cusp at $c$.) The opposite-side cases with two and with three planes are depicted schematically in aff( $[o, p, q]$ ) in Figures 6.10(a), and (b), respectively.

In the same-side case, $L_{1}$ and $L_{2}$ are not sufficient even within the circumcylinder. Instead, we define $L_{c}$ to be the vertical plane that contains $r$ and $p^{*}$, and $L_{f}$ to be the one that contains $f$ and $r$. Thus $L_{f}$ corresponds to either $L_{1}$ or $L_{2}$ in the opposite-side case, but $L_{c}$ is necessarily more restrictive, isolating the close vertex from the allowable region,
set $r$ to midpoint of $\overparen{q O}$
set $c, f$ to whichever of $o$ and $q$ is closer, farther to,from $p^{*}$
if $o$ and $q$ are on opposite sides of $\left[p, p^{*}\right]$ then
set $L_{1}$ to vertical constraint plane through $o$ and $r$
set $L_{2}$ to vertical constraint plane through $q$ and $r$
if $\left|\left[c, p^{*}\right]\right|<\left|\left[p^{*}, r\right]\right|$ then
set $a$ to midpoint of $\overparen{r c}$
set $L_{t}$ to vertical constraint plane through $c$ and $a$
end if
else
set $L_{c}$ to vertical constraint plane through $p^{*}$ and $r$
set $L_{f}$ to vertical constraint plane through $f$ and $r$
set $z$ to $p^{*}+\frac{\overrightarrow{f p^{*}}}{\left\|\overrightarrow{f p^{*}}\right\|}+\tan \tau\left(\frac{\overrightarrow{p f}}{\|\overrightarrow{p f}\|} \times \frac{\overrightarrow{p c}}{\|\overrightarrow{p c}\|}\right) \times \frac{\overrightarrow{f p^{*}}}{\left\|\overrightarrow{f p^{*}}\right\|}$
set $L_{t}$ to vertical constraint plane through $c$ and $z$
end if
Algorithm 3: Incident edge constraints
rather than going through it. We also always need a plane $L_{t}$, tangent to the boundary of the feasible region at $p^{*}$. In this case $L_{t}$ is the vertical plane through $p^{*}$ such that it makes an angle $\tau$ with $\left[f, p^{*}\right]$. The angle $\tau$ is defined by Equation (D.10):

$$
\tan \tau=\frac{k \sin \gamma}{(1+k \cos \gamma)}
$$

where $\gamma=\angle o p q$, and $k=\frac{\|\left[f, p^{*}\right] \mid}{\|\left[c, p^{*}\right] \mid}$. The same-side case is depicted in Figure 6.10(c).
The construction of the constraint planes for incident edges is summarized in Algorithm 3. The derivation and demonstration of the correctness of these constraints is detailed in Appendix D.

Remark 6.7 In the original implementation of the decimation algorithm [DZM07a], only planes $L_{1}$ and $L_{2}$ were used for all cases. The output was a self-Delaunay mesh on all models tested, indicating that the extra constraints imposed by $L_{c}$ and $L_{t}$ are not generally needed in practice. The output of the corrected algorithm is qualitatively similar to that of the original, but in some cases the Hausdorff error is slightly larger in the output of the corrected version.

Allowable region for an edge collapse: Referring to Figure 6.7, for an edge collapse $[u, v] \rightarrow w$, the allowable region for $w$ is the intersection of all the allowable regions defined for the incident and subtending edges for $w$. If this set is empty, then $[u, v]$ will not be collapsed.

### 6.4 Experimental results

We have tested our mesh edge flipping and refinement algorithm (Algorithm 2) on a few dozen mesh models and statistics collected on some representative data are reported in Table 6.1. The first group of models are well-known and serve as examples of typical datasets. Note that the Stanford bunny and Max Plank models have boundary edges. In the second group, we choose two low-resolution meshes that have thin structures; they were obtained via QSlim [GH97]. These models have a greater percentage of 3- or 2-exposed tets compared with those from the first group and provide a more rigorous test for the refinement component of our algorithm. The final group contains two meshes that were produced by the remeshing algorithm of Peyré and Cohen [PC03]; these meshes are generally quite close to being self-Delaunay.

Table 6.1 shows that our algorithm consistently terminates after flipping and splitting a small fraction of the mesh edges. Denser models tend to incur smaller percentages of such operations. The Peyré hand and horse models, both coarsely sampled, were obtained from approximate geodesic Delaunay triangulations of points sampled on the original fine mesh surfaces. In other words, these meshes approximate iDt-meshes of the smooth surfaces represented by the original high resolution models. However, the meshes produced do have a small fraction of edges that are nlD.

With respect to angle quality, although the smallest angle is not required by theory to increase via Delaunay remeshing, Table 6.1 shows that in practice it generally does with few exceptions. In no case did the size of the maximum angle, or the percentage of either small $\left(<30^{\circ}\right)$ or large $\left(>120^{\circ}\right)$ angles increase in the self-Delaunay meshes.

In the last column of Table 6.1, we report the approximation error $\epsilon$ of the self-Delaunay meshes produced, measured using the well-known Metro tool [CRS98]. As can be seen, the geometric approximation error tends to be large for coarse models. In contrast, errors associated with the self-Delaunay versions of the corresponding full-resolution models are much smaller; see results for models from the first group.

| Mesh | \#E | Flips(\# \| \%) |  | Splits | Min angle |  | Max angle |  | \% Small |  | \% Large |  | Error ( $\epsilon$ ) |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Horse | 59547 | 2957 | (5.0\%) | 28 | 1.7 | 1.6 | 171.3 | 164.3 | 8.7 | 7.9 | 1.3 | 1.2 | 0.4741\% |
| Hand | 74997 | 10935 | (14.6\%) | 0 | 2.1 | 3.2 | 175.3 | 156.9 | 15.5 | 8.5 | 3.3 | 0.6 | 0.1061\% |
| Bunny | 104288 | 2202 | (2.1\%) | 82 | 0.5 | 7.8 | 177.6 | 146.1 | 2.3 | 2.1 | 0.2 | 0.0 | 0.2367\% |
| Igea | 165000 | 19349 | (11.7\%) | 0 | 0.1 | 6.1 | 179.8 | 148.3 | 10.0 | 5.2 | 2.2 | 0.3 | 0.1037\% |
| Isis | 562926 | 13497 | (2.4\%) | 2 | 0.4 | 3.1 | 176.8 | 163.6 | 2.5 | 2.4 | 0.1 | 0.0 | 0.1591\% |
| Max Planck | 597211 | 12157 | (2.0\%) | 12 | 0.1 | 5.1 | 178.9 | 145.2 | 1.7 | 1.1 | 0.3 | 0.0 | 0.0992\% |
| Coarse hand | 294 | 99 | (33.7\%) | 34 | 4.7 | 12.4 | 167.4 | 142.1 | 27.2 | 11.4 | 7.8 | 1.5 | 4.3245\% |
| Coarse horse | 1050 | 231 | (22.0\%) | 36 | 3.5 | 12.9 | 171.3 | 133.5 | 17.9 | 6.9 | 5.5 | 0.3 | 1.6475\% |
| Peyre hand | 894 | 22 | (2.5\%) | 0 | 24.7 | 31.0 | 129.9 | 112.5 | 0.6 | 0.0 | 0.2 | 0.0 | 1.4095\% |
| Peyre horse | 1944 | 16 | (0.8\%) | 4 | 12.0 | 19.4 | 136.4 | 113.8 | 0.2 | 0.2 | 0.1 | 0.0 | 0.8684\% |

Table 6.1: Output statistics for our edge flipping and refinement algorithm (Algorithm 2) on several mesh models. Input mesh sizes are measured by $\# E$, the edge count. We report the number of edge flips performed until termination both in absolute number and as a percentage of $\# E$, as well as the number of edge splits due to unflippable and boundary edges. The next four double columns have before (left sub-column) and after (right sub-column) figures for the minimum (Min angle) and maximum (Max angle) face angles (in degrees) in the meshes, as well as the percentage of angles that are less than $30^{\circ}$ (\% Small) and the percentage of angles that exceed $120^{\circ}$ (\% Large). The last column reports the approximation error given by Metro [CRS98], as a percentage of the Hausdorff distance between the meshes against the length of the bounding box diagonal. The longest running time is recorded on the Max Plank, which took 12.2 seconds to process on a 2.4 GHz Opteron processor.

When used on its own, the geometry-preserving refinement scheme may produce excessive and poorly distributed vertices. However, when strict geometry preservation is relaxed and we allow flipping of edges that have small dihedral angles, we obtain an algorithm that produces more pleasing results. Essentially this is edge flipping with feature preservation, where physical mesh edges with dihedral angles exceeding a user-set threshold are flagged and they are split according to the scheme of Section 6.1 as they appear in the priority queue. A similar argument to that of Section 6.1.2 carries the termination guarantee to this modified algorithm.

(a) Original: $\# V=6,002$.

(b) $\# V=6,002$;
$\epsilon=0.385 \%$.

(c) $\# V=15,560$;
$\epsilon=0.000 \%$.

(d) $\# V=7,509$;
$\epsilon=0.071 \%$.

Figure 6.11: A close-up of results from Delaunay remeshing on the dolphin model. Mesh vertex count $(\# V)$ and Metro error $(\epsilon)$, measured against the original model, are given below the figures. (a) Original. (b) After unconstrained edge flipping. (c) After geometrypreserving refinement. (d) After feature-preserving edge flipping with a threshold of $10^{\circ}$ on dihedral angles.

Figure 6.11 (a) shows a close-up of a coarsely sampled dolphin model ( 6,002 vertices). The mesh edge flipping algorithm does a poor job of preserving the detailed features in the initial shape, as shown in (b), while geometry-preserving refinement creates a large vertex count ( 15,560 vertices) and uneven vertex distributions, as shown in (c). However, edge flipping with feature preservation presents a nice compromise, shown in (d), where the vertex count is reduced to 7,509 and the Metro error $\epsilon$ is reduced from $0.385 \%$ to $0.071 \%$, when a dihedral angle threshold of $10^{\circ}$ is used.

A tradeoff between feature preservation and vertex count verses the threshold dihedral angle is illustrated in Figure 6.12. This graph is produced from the Isis model of 187,644 vertices. If the threshold is too high, many vertices are introduced in order to preserve geometry that is probably more associated with the initial discretization than with the intended model. For coarser initial models, the graphs have a similar appearance, but the


Figure 6.12: Graph of relative increase in vertex count verses the threshold dihedral angle for feature preserving edge flipping of the Isis model (187,644 vertices initially).
scale on the vertical axis is much larger. For example, the same graph generated from an initial Isis model of 600 vertices (QSlimmed) has a scale an order of magnitude larger on the vertical axis. Thus the cost, in terms of vertex count, of geometric fidelity, is much greater (relatively) for a coarsely sampled model.

Finally, a multiresolution family of meshes produced by our Delaunay mesh decimation algorithm is given in Figure 6.13. The decimation generally performs well in preserving features in its attempts to minimize the quadric errors. However, our current implementation does not employ lazy evaluation or other heuristics to speed up the optimization and is thus very slow.


Figure 6.13: Hand model at three resolutions produced by our decimation algorithm. The Metro errors are: $0.011 \%$ for (a), $0.159 \%$ for (b) and $1.372 \%$ for (c), all measured against a Delaunay remeshing obtained from the original model.

### 6.5 Discussion

We have seen that the Delaunay edge flipping algorithm, well known in the context of planar triangulations, can be extended to a Delaunay extrinsic edge flipping algorithm on general manifold triangle meshes. However, there are two new problems which arise in this context. One is that extrinsic edge flipping causes geometric distortion, and the artifacts can be severe if edges with large dihedral angles are flipped, as shown in Figure 6.14. The second, and more fundamental, problem is that of unflippable edges. It may not be possible to produce a manifold self-Delaunay mesh by edge flipping alone.

Our solution to both these problems is the geometry preserving Delaunay refinement algorithm. This algorithm is of little practical interest as a stand alone algorithm. We do not have a theoretical bound on the number of vertices that may be inserted, but even if the expected $\mathcal{O}(n)$ theoretical bound is established, it will be of little consolation. As exemplified


Figure 6.14: The original model (left), shows serious artifacts after Delaunay extrinsic edge flipping (right).
in Figure 6.11(c), the number of vertices inserted can exceed the number of vertices in the original model. However, the algorithm does serve to supplement the Delaunay extrinsic edge flipping algorithm to ensure that a final self-Delaunay mesh is reached. In this context, experiments indicate that the number of vertices inserted is small. Also, when used in tandem, edge flipping and refinement can complement each other to minimize geometric distortion while keeping vertex insertion at acceptable levels.

Thus the feature preserving edge flipping with refinement algorithm can be tuned to the appropriate balance between geometric distortion on the one hand, and vertex count on the other hand. Near the extreme that is geometry preserving refinement, besides the high vertex count, the distribution of the vertices is irregular and reflects more the characteristics of the initial triangulation, than the curvature inherent in the represented geometry. However, these self-Delaunay meshes make ideal input to our Delaunay decimation algorithm whose output has a nice triangle quality and vertex distribution.

## Chapter 7

## Analysis of Delaunay extrinsic edge flips

In this chapter we make further investigations into the properties of the Delaunay extrinsic edge flipping algorithm. The algorithm is of particular interest because it is a convenient tool for gaining insight into self-Delaunay meshes themselves. Specifically, we are interested in characterizing sampling conditions which would ensure that a point set $P$ sampled from $S$ admits a smooth self-Delaunay mesh with $P$ as its vertex set. A natural approach to this problem is to identify conditions under which the Delaunay extrinsic edge flipping algorithm will not encounter an unflippable nlD edge. Then the existence of a smooth self-Delaunay mesh on $P$ would be guaranteed by sampling criteria which ensure that a known mesh structure, such as the rDt, meets these conditions.

In the geodesic Delaunay edge flipping algorithm described in other works [ILTC01, BS07, Gli05, FSBS06], unflippable edges are not an issue because the triangulation need not be proper. Given a mesh $M$, its extrinsic triangulation is proper, and may be taken as the input for an edge flipping algorithm. The geodesic Delaunay edge flipping algorithm, and the extrinsic Delaunay edge flipping algorithm both employ the same criteria to decide whether an edge should be flipped: the edge must be nlD. Edge $e$ will be unflippable for the Delaunay extrinsic edge flipping algorithm if and only if the geodesic edge flipping algorithm would yield a non-proper triangulation by flipping $e$.

The conditions under which geodesic Delaunay edge flipping maintains a proper triangulation have not been identified. We know that if the iVd of $M$ is well formed, then the
iDt will be proper. It is also known that, assuming a sufficient vertex density on $\mathcal{M}$, any proper triangulation may be flipped to any other proper triangulation by a sequence of edge flips such that all intermediate triangulations are proper [Neg94, Kin02]. However, this latter result from combinatorial edge flipping does not inform us as to whether any or all sequences of geodesic Delaunay edge flips will maintain a proper triangulation.

Extrinsic edge flipping algorithms for optimizing various curvature-based functionals on meshes have been studied, for example in [vDA95, DHKL01, ABR06], but the question of unflippable edges has not been broached in any detail in these works. An explanation for this is that it is recognized that the edge flipping algorithms are not guaranteed to converge to global optima for the considered functionals, regardless of whether or not unflippable edges are encountered. Thus in that context unflippable edges are a detail that is overshadowed by a larger issue.

Aichholzer et al. [AAH02] study extrinsic edge flipping in a general context and show that, if $M$ and $M^{\prime}$ are distinct genus zero meshes on a common vertex set $P$, then, even when $P$ is in convex position, it is in general not possible to transform $M$ into $M^{\prime}$ by extrinsic edge flips such that the intermediate meshes are all embedded. In a subsequent work, Alboul [Alb03], made a distinction between local and global self-intersections. In our terminology a local self-intersection implies a singular mesh, whereas a global selfintersection refers to a manifold triangle mesh that is not embedded. In the context of the optimization studied in [Alb03], it is the global self-intersections that cause a problem, and singular meshes are accommodated by the algorithm.

We find the opposite situation for the case of Delaunay extrinsic edge flips. We show in Section 7.1 that if $M$ is a smooth mesh, then it has no unflippable nlD edges. The demonstration depends critically on the fact that a smooth mesh is a manifold triangle mesh. This result highlights the importance of smooth meshes in the context of Delaunay extrinsic edge flipping, although we demonstrate that a slightly weaker condition, which we call non-sharp, suffices.

However, in order to use this result to guarantee that no unflippable nlD edges are encountered, we must ensure that the Delaunay extrinsic edge flipping algorithm will not produce a mesh that is not smooth. This motivates our study, in Section 7.2, of the behaviour of the triangle circumradii under Delaunay extrinsic edge flips. We find that there are nlD hinges in which a Delaunay edge flip will increase the largest circumradius, and such hinges cannot be avoided by any density-based sampling criteria.

In Section 7.3 we explore a different aspect of the existence question. We ask whether a self-Delaunay mesh must exist as a substructure of $\mathcal{D}^{3}(P)$. We show that a cocone mesh may have an nlD edge whose flip-tet does not belong to $\mathcal{D}^{3}(P)$. This implies that the mesh resulting from flipping that edge will not be a substructure of $\mathcal{D}^{3}(P)$. However, the construction does not provide an obstruction to a self-Delaunay mesh belonging to $\mathcal{D}^{3}(P)$, because we have no example where such a flip results in a self-Delaunay mesh.

### 7.1 Non-sharp and smooth meshes

Our focus is on meshes that are models for $S$, a smooth surface embedded in $\mathbb{R}^{3}$. In that context it is natural to consider smooth meshes. We now consider the problem of unflippable edges encountered by the Delaunay extrinsic edge flipping algorithm, and we find another motivation for interest in smooth meshes: within the domain of smooth meshes, unflippable nlD edges cannot occur.


Figure 7.1: Non-sharp, but not smooth.

The obstruction to unflippable nlD edges is a property of smooth umbrellas which warrants an independent definition. The umbrella of a vertex $p \in M$ is non-sharp if: For any face $t$ incident to $p$, any line $\ell$ which intersects the interior of $t$ perpendicularly does not intersect another face in the umbrella of $p$. Define the orthogonal shadow of a face $t$ as the triangular prism, $\operatorname{sh}(t)$, consisting of those points in $\mathbb{R}^{3}$ which project orthogonally onto $t$. Then the non-sharp condition demands that no interior point in a triangle in the umbrella of a vertex of $t$ can lie inside $\operatorname{sh}(t)$.

We say that vertex $p$ is non-sharp if has a non-sharp umbrella. An edge $e$ in $M$ is non-sharp if the angle between consistent normals of the faces incident to $e$ is less than $\pi / 2$. Any edge that is incident to a non-sharp vertex must itself be non-sharp. A mesh is non-sharp if all its vertices are non-sharp. If a vertex or edge is not non-sharp, then it is sharp.

Lemma 7.1 A smooth umbrella is non-sharp.
Proof Suppose $p$ has a smooth umbrella, and suppose that faces $f$ and $t$ have $p$ as a vertex and that an interior point $x \in f$ intersects $\operatorname{sh}(t)$. Since $\operatorname{sh}(t)$ is convex, if $x$ lies inside $\operatorname{sh}(t)$, then so will any point on the line segment $[x, p] \subset f$. Consider the Euclidean ball $B=B_{\mathbb{R}^{3}}(p ; r)$ centred at $p$ and with radius $r$ small enough that the one ring of $p$ is exterior to $B$. Since $U(p)$ is manifold, it separates $B$ into two pieces, $B^{+}$and $B^{-}$and we choose an orientation on the face normals such that they point towards $B^{+}$.

Now consider a line $\ell$ perpendicular to $t$ which intersects $t$ and $f$ inside $B$. Then $\hat{\boldsymbol{n}}_{t}$ generates $\ell$. Since $U(p)$ is smooth,
 every other face normal has a positive scalar product with $\hat{\boldsymbol{n}}_{t}$. Thus when traversing $\ell$ in the $\hat{\boldsymbol{n}}_{t}$ direction, we will necessarily pass through $t$ from $B^{-}$to $B^{+}$, and likewise when we pass through $f$. But this is impossible unless we pass from $B^{+}$to $B^{-}$through a third face between $t$ and $f$, and this would violate the assumption that $U(p)$ is smooth.

Thus all smooth meshes are non-sharp. However the converse is not true. Figure 7.1 shows an example where two triangles in the umbrella of $p$ which lie on opposite sides of $p$ have a relative transverse twist that exceeds $\pi / 2$, and yet no line perpendicular to one will pierce the other.

The non-sharp criterion is weaker than the smooth criterion, but if $p$ is a vertex of valence three, then the non-sharp property is equivalent to the smooth property, and it is satisfied provided the incident edges are all non-sharp.

The following lemma captures the important connection between Delaunay extrinsic edge flipping and the non-sharp property. Roughly speaking, it says that the opposing edge to a non-sharp nlD edge $e$ lies "underneath" the hinge associated with $e$. It can be compared to the observation that in the plane, the opposing edge to an nlD edge lies within the flip quad (because the flip quad is convex), thus guaranteeing the validity of planar Delaunay edge flipping.

Lemma 7.2 If $e$ is a non-sharp nlD edge, then the opposing edge, $e^{\prime}$, intersects the interior of the orthogonal shadow of each of the faces in the hinge of $e$.

Proof Let $(u, e, v)$ be a non-sharp hinge on $e=[p, q]$ with opposing edge $e^{\prime}$. Let $t_{1}=[u, p, q]$, and let $v^{\prime}$ be the orthogonal projection of $v$ onto aff $\left(t_{1}\right)$. Since $\angle u p q+\angle q p v<\pi$, it follows that $\angle u p q+\angle q p v^{\prime}<\pi$; the vector $\vec{p}$ has a positive component in the $\hat{\boldsymbol{n}}_{t_{1}} \times \overrightarrow{u p}$ direction. Likewise $\angle u q p+\angle p q v^{\prime}<\pi$. Therefore
 $v^{\prime}$ must lie in the wedge defined by $\angle q v p$. Since $e$ is non-sharp, $v^{\prime}$ must lie outside of $e$, and so the segment $\left[v^{\prime}, u\right]$ must cross the interior of $e$. Since $\left[v^{\prime}, u\right]$ is the image of $e^{\prime}$, it follows that $e^{\prime}$ intersects the interior of $\operatorname{sh}\left(t_{1}\right)$. In the same way, we see that $e^{\prime}$ must intersect the interior of the orthogonal shadow of $t_{2}=[v, q, p]$.

Lemma 7.2 immediately implies the following observation, which asserts that the flip-tet of an nlD edge has a flat appearance.

Lemma 7.3 In the flip-tet of a non-sharp nlD edge, all edges which are shared by a pre-flip and a post-flip triangle are sharp.

Lemma 7.2 leads to the main observation of this section, which is that an nlD edge in a non-sharp mesh can always be flipped.

Lemma 7.4 In a non-sharp mesh there are no unflippable nlD edges.
Proof Suppose $e=[p, q]$ is an unflippable nlD edge in a non-sharp mesh. Then its opposing edge $e^{\prime}=[u, v]$ already belongs to the umbrella at $u$. By Lemma 7.2, $e^{\prime}$ intersects the interior of the orthogonal shadow of the triangle $t=[u, p, q]$. But $t$ also belongs to the umbrella at $u$, contradicting the hypothesis that the mesh is non-sharp.

Also, the non-sharp property is sufficient to ensure that no Delaunay edge flip will cause the new mesh to become singular.

Proposition 7.5 If $M$ is a non-sharp manifold triangle mesh, then the mesh $M^{\prime}$ obtained by performing a Delaunay extrinsic edge flip on $M$ will also be a manifold triangle mesh.

Proof Suppose $e=[p, q]$ is nlD with opposing edge $e^{\prime}=[u, v]$. Since, by Lemma 7.4, $e$ is flippable, it is sufficient to show that none of the other triangles in the umbrellas at $p, q, u$ and $v$ will intersect the new faces $t_{1}^{\prime}=[p, v, u]$ or $t_{2}^{\prime}=[q, u, v]$.

To see this, observe that the proof of Lemma 7.2 reveals that any point $x \in e^{\prime}$ must lie in the orthogonal shadow of either $t_{1}=[u, p, q]$ or $t_{2}=[v, q, p]$. From this it follows by the
convexity of the orthogonal shadows, that any point in $t_{1}^{\prime}$ or $t_{2}^{\prime}$ must also lie in $\operatorname{sh}\left(t_{1}\right)$ or $\operatorname{sh}\left(t_{2}\right)$. Indeed, if $x \in t_{1}^{\prime}$, then the line through $p$ and $x$ intersects $e^{\prime}$ in a point $y$ that is in at least one of $\operatorname{sh}\left(t_{1}\right)$ and $\operatorname{sh}\left(t_{2}\right)$. Since $p$ lies in both shadows, $x$ must also lie in the same orthogonal shadow as $y$.

Thus if a face $f$ intersects $t_{1}^{\prime}$ or $t_{2}^{\prime}$ it cannot belong to an umbrella of a vertex of the hinge of $e$ without violating the non-sharp condition on $M$.

Although our focus is on Delaunay extrinsic edge flips, the essential property of an nlD hinge that was exploited in the proof of Lemma 7.2, is that it unfolds to a convex quadrilateral. Thus the observation of Proposition 7.5 can be generalized:

Corollary 7.6 If $e$ is an edge in a non-sharp manifold mesh $M$, and if the hinge on $e$ is isometric to a convex planar quadrilateral, then the mesh obtained by flipping $e$ is a manifold triangle mesh.

A necessary and sufficient condition for a planar edge flipping algorithm to maintain a valid mesh is that it only flips edges whose flip-quad is convex. Corollary 7.6 reveals that in the case of a smooth triangle mesh this same condition ensures validity, so long as the mesh remains smooth.

### 7.2 Delaunay extrinsic edge flips and triangle circumradius

If a mesh is smooth, then it has no unflippable nlD edges. However, the Delaunay extrinsic edge flipping algorithm may encounter unflippable edges even if the initial mesh is smooth. Without additional constraints on the input mesh, there is no guarantee that the algorithm will maintain mesh smoothness, as demonstrated in Figure 7.2 . A desire to characterize the needed additional constraints is the motivation behind the work presented in this section.

As discussed in Section 2.4, smoothness will be ensured if we are able to provide a sufficient bound on the triangle circumradii. In this section we examine the relative circumradii of triangles in the flip-tet of an nlD edge. We rely heavily upon Lemma 4.14, which for convenience we restate here for the specific case of triangles:

Lemma 7.7 Suppose triangles $t$ and $t^{\prime}$ share an edge $e$ and that $t^{\prime}$ is contained in $B_{t}$. If $e$ subtends an acute angle in $t^{\prime}$, then $r_{t}>r_{t^{\prime}}$.


Figure 7.2: From a smooth mesh we obtain an unflippable nlD edge. (a) The initial mesh is smooth (the boundary is irrelevant; we could cap this object with smooth self-Delaunay domes). (b) There is a solitary vertex in the interior of the cylinder. (c) The Voronoi diagram of the mesh vertices is not well formed (We are looking at the back; the solitary vertex is on the other side). (d) The Delaunay edge flipping algorithm eventually encounters an unflippable edge.

In what follows we will consider an nlD edge $e=[p, q]$ with opposing edge $e^{\prime}=[u, v]$ as shown in the figure. Let the pre-flip triangles be $t_{1}=[u, p, q]$ and $t_{2}=[p, v, q]$ and without loss of generality, assume $r_{t_{1}} \geq r_{t_{2}}$. The post-flip triangles will be denoted $t_{1}^{\prime}$ and $t_{2}^{\prime}$ and such that $r_{t_{1}^{\prime}} \geq r_{t_{2}^{\prime}}$.

In the case of a planar triangulation, a Delaunay edge flip will always


Flip-tet yield $r_{t_{1}^{\prime}}<r_{t_{1}}$ and $r_{t_{2}^{\prime}}<r_{t_{2}}$, however for Delaunay extrinsic edge flips in a non-planar mesh, the first inequality does not always hold. We do at least maintain the second inequality for all flips:

Lemma 7.8 In an Delaunay extrinsic edge flip $\left(t_{1}, t_{2}\right) \rightarrow\left(t_{1}^{\prime}, t_{2}^{\prime}\right)$ on a non-sharp mesh, the smaller of the circumradii of the post-flip triangles is smaller than the circumradius of each of the pre-flip triangles. I.e., $r_{t_{2}^{\prime}}<r_{t_{2}}$.

Proof By Lemma 4.11, and Lemma 4.10, $t_{2}$ has an acute angle subtended by $e$, and it does not yield a Gabriel certificate to $e$. We will show that Lemma 7.7 can be used to compare at least one of $r_{t_{1}^{\prime}}, r_{t_{2}^{\prime}}$ with $r_{t_{2}}$. To that end, we need to verify that either $t_{1}^{\prime}$ or $t_{2}^{\prime}$ has an acute angle at $u$.

If $\angle p u q$ is acute, then it follows from Lemma 7.2 that both $\angle p u v$ and $\angle v u q$ are acute. Indeed, the scalar product of $\overrightarrow{u v}$ and $\overrightarrow{u p}$ is determined by the component of $\overrightarrow{u v}$ that lies in the plane of $[u, p, q]$, and will clearly be positive. Thus if $\angle p u q$ is acute the result follows from Lemma 7.7 by comparing $t_{2}$ and $t_{2}^{\prime}$ on their common edge.

If $\angle p u q$ is not acute, Lemma B. 8 implies that at least one of the angles $\angle p u v$ and $\angle v u q$ must be acute. This follows because Lemma 7.3 ensures that $u$ is a sharp vertex in the flip-tet and the Gauss map transforms the valence three vertex $u$ into a spherical triangle whose angles are the supplements of the corresponding face angles at $u$. Therefore we can again apply Lemma 7.7 to compare $r_{t_{2}}$ with at least one of $r_{t_{1}^{\prime}}$ and $r_{t_{2}^{\prime}}$, and the result follows.

Although the smaller circumradius will always be reduced, the behaviour of the larger circumradius is not so easily tamed. In the following we classify the flip-tets according to their behaviour in this regard. Since Lemma 7.7 is the primary tool used to compare circumradii, the Gabriel properties of the flip-tet become a natural way to classify the different cases. This classification leads to three possible cases.

The first case, when the initial nlD edge has no Gabriel certificates, can be handled analogously to the planar case, as demonstrated in Cheng and Dey [CD07][Lemma 3.2]:

Lemma 7.9 If $e$ has no Gabriel certificates, then the largest post-flip circumradius will be smaller than the largest pre-flip circumradius. I.e., $r_{t_{1}^{\prime}}<r_{t_{1}}$.

Proof Since $t_{1}^{\prime}$ shares an edge with $t_{1}$ and another with $t_{2}$, one of these edges must subtend an acute angle in $t_{1}^{\prime}$. Since neither $t_{1}$, nor $t_{2}$ yields a Gabriel certificate to $e$, we can apply Lemma 7.7 to get the required comparison with at least one of $r_{t_{1}}$ and $r_{t_{2}}$.

The remaining two cases arise when $e$ has a single Gabriel certificate. If the opposing hinge has only a single Gabriel certificate, then the edge flip results in circumradius growth.

Lemma 7.10 If nlD edge $e$ is flipped to an opposing edge $e^{\prime}$ which has only a single Gabriel certificate, then the largest post-flip circumradius will exceed both of the pre-flip circumradii. I.e., $r_{t_{1}^{\prime}}>r_{t_{1}}$.

Proof Lemma 4.11 and Lemma 4.10 imply that $t_{1}^{\prime}$ has an acute angle subtended by $e^{\prime}$, and does not yield a Gabriel certificate to $e^{\prime}$. Also, since $t_{1}$ and $t_{1}^{\prime}$ share an edge that terminates in $u$, and $t_{1}$ has an obtuse angle at $u$ (by Lemma 4.11), it follows that the common edge must subtend an acute angle in $t_{1}$. Therefore the result follows from Lemma 7.7.

The only remaining case to consider is when an edge with a single Gabriel certificate gets flipped to one with two Gabriel certificates. In this case further information is needed to determine whether circumradius growth occurs.

Lemma 7.11 Suppose edge $e$ has a single Gabriel certificate and it is flipped to an edge $e^{\prime}$ that has two Gabriel certificates. If $t_{1}^{\prime}$ has an acute angle at $u$, then $r_{t_{1}^{\prime}}<r_{t_{1}}$. Otherwise, $r_{t_{1}^{\prime}}<r_{t_{1}}$ iff the acute angle in $t_{1}^{\prime}$ at $v$ is larger than the angle in $t_{1}$ that is subtended by the edge common to $t_{1}$ and $t_{1}^{\prime}$.

Proof Lemma 4.11 together with Lemma 4.10 implies that $t_{2}$ does not yield a Gabriel certificate to $e$. Therefore, if $t_{1}^{\prime}$ has an acute angle at $u$ then Lemma 7.7 ensures that $r_{t_{1}^{\prime}}<r_{t_{2}}$.

If $t_{1}^{\prime}$ does not have an acute angle at $u$, then the final statement follows from the fact that for a triangle on a fixed edge subtending an acute angle, the circumradius grows as the angle decreases.

In Section 3.2.1, we observed that most of the functionals which are optimized by the Delaunay edge flip algorithm in the plane, or on a fixed pwf surface, are not optimized by the Delaunay extrinsic edge flip algorithm. Lemma's 7.10 and 7.11 confirm that the triangle circumradius is no exception to this trend.

If we could guarantee that triangle circumradius would not increase as a result of a Delaunay extrinsic edge flip, then mesh


Recall Figure 4.4 smoothness would be ensured if the sampling density were constant (i.e. not adapted to the local feature size). However, even if a local uniformity constraint is imposed on the sample
distribution, nlD hinges satisfying the hypothesis of Lemma 7.10 may still occur. Indeed, the same flip-tet, depicted in Figure 4.4, that provided an obstruction to the existence of closed Gabriel meshes, may be furnished as an example. The observations of Section D. 1 indicate that we may ensure that edge $[p, q]$ is nlD, while the ratio of the circumradius to shortest edge in $\sigma$ is arbitrarily close to unity.

If a smooth self-Delaunay mesh on $P$ exists, then we expect that the circumradii of the triangles will be similar in size to that of nearby triangles in the rDt. However, the observations of this section imply that a bound on the triangle circumradii cannot be obtained through the flip algorithm unless there is a definite bound on the number of flips that may occur on edges incident to a given vertex.

### 7.3 Relation to the ambient Delaunay tetrahedralization

We have seen in Section 3.3.2 that the rDt of $P \subset S$ is not a self-Delaunay mesh in general. However, the rDt belongs to the family of cocone meshes on $P$ which may include many distinct meshes, each of which is a good representation of $S$ and a substructure of $\mathcal{D}^{3}(P)$. So we may ask the question: do the cocone meshes on $P$ include a self-Delaunay mesh?

The question is intimately linked to the role of sliver tetrahedra, as elements in $\mathcal{D}^{3}(P)$ and as flip-tets. Sliver tetrahedra are loosely defined as those tetrahedra whose vertices all lie near a common great circle in the circumsphere. A precise or quantifiable definition is rarely given, but to be specific we could say that $\sigma$ is a sliver tetrahedron if $d_{\mathbb{R}^{3}}\left(c_{\sigma}, c_{t}\right)<\epsilon r_{\sigma}$ for any triangle $t \subset \sigma$. The parameter $\epsilon$ would depend on the corresponding parameter governing the sampling radius with respect to the local feature size.

In a sliver tetrahedron the circumradius of each triangle face is almost as big as the circumradius of the tetrahedron itself, and this would be another way of characterizing a sliver, but for the purposes of the current discussion, the separation of circumcentres is the more enlightening criterion. In Section 3.3.2 we constructed an nlD hinge in an rDt. The associated flip-tet, $\sigma$, is a sliver because $c_{\sigma}$ is bounded between $S$ and the hinge. In fact, the flip-tet of any nlD edge in a cocone mesh must be a sliver tetrahedron.

To see this, observe that if $\left(t_{1}, t_{2}\right)$ is nlD, then $c_{\sigma}$ must lie on the convex side of the hinge (it must be in the exterior half-space of at least one of the triangles). Considering $\sigma$ in isolation, the Voronoi edges of $t_{1}$ and $t_{2}$ extend from $c_{\sigma}$, away from the hinge. Since $t_{1}$ and $t_{2}$ are cocone triangles, these Voronoi edges must intersect the cocone regions of the
corresponding vertices of the hinge. Since the cocone regions lie close to $S$, and also the distance from any point $x$ in the hinge to $S$ is $\mathcal{O}\left(\epsilon^{2} \rho_{f}\left(\xi_{S}(x)\right)\right)$, where $\epsilon \rho_{f}(x)$ is the sampling radius, it follows that $c_{\sigma}$ must lie close to the hinge. The true Voronoi edges dual to $t_{1}$ and $t_{2}$ in the Voronoi diagram of $P \subset \mathbb{R}^{3}$ will be subsets of the Voronoi edges obtained when $\sigma$ is considered in isolation, thus considering the true Voronoi edges will not relax the outcome.

Thus the flip-tets of nlD edges in a cocone mesh are slivers. Also, it is due to the existence of sliver tetrahedra in $\mathcal{D}^{3}(P)$ that the family of cocone meshes include meshes other than the rDt [Dey07]. We could hope that a Delaunay edge flip in a cocone mesh always results in another cocone mesh. A necessary condition for this to be the case would be that the flip-tet of each nlD edge belongs to $\mathcal{D}^{3}(P)$. However, we show here that this is not the case.


Figure 7.3: (a) The nlD hinge $(u, e, v)$ contains faces from two distinct tetrahedra in $\mathcal{D}^{3}(P)$ : $\sigma^{\prime}=[p, u, w, q]$ is shown in green, and $\sigma^{\prime \prime}=[p, w, q, v]$ is shown in blue, and their shared face, $t=[p, w, q]$ is shown in red. The five vertices are nearly coplanar, with $t$ lying in the plane of the page, $u$ lying just above it, and $v$ below it. The diametric sphere of right triangle $t_{1}=[p, u, q]$ is depicted by the black circle. The other two vertices lie inside it. (b) A vertical cross section perpendicular to $e$, depicting the relative locations of circumcentres. Triangle $t_{2}=[v, p, q]$ is depicted as the thick solid blue line. The thin solid blue line perpendicular to this runs through $c_{t_{2}}$ and $c_{\sigma^{\prime \prime}}$. Similarly for the triangles $t$ and $t_{1}$.

We construct a hinge ( $u, e, v$ ) on nlD edge $e=[p, q]$ such that the flip-tet $\sigma=[u, q, p, v]$ does not belong to $\mathcal{D}^{3}(P)$. We introduce a fifth vertex, $w$, and show that the tetrahedra
$\sigma^{\prime}=[p, u, w, q]$ and $\sigma^{\prime \prime}=[p, w, q, v]$ share the face $t=[p, w, q]$ such that $t$ is locally Delaunay. See Figure 7.3(a). In other words, $S_{\sigma^{\prime}}$ does not contain $v$, and $S_{\sigma^{\prime \prime}}$ does not contain $u$.

As described in Figure 7.3, we make $t_{1}=[p, u, q]$ a right triangle, and we push $v$ inside $S_{t_{1}}$ such that $c_{t_{2}}$ is a distance $\delta_{2}$ from $c_{e}=c_{t_{1}}$, where $t_{2}=[v, p, q]$. This ensures that $t_{2}$ is obtuse and thus ( $u, e, v$ ) is an nlD hinge. We let $t$ lie in the plane of the page, and $w$ is pushed inside $S_{t_{1}}$ such that $c_{t}$ is a distance $\delta$ from $c_{e}$. Vertex $u$ is lifted above the page so that $t_{1}$ makes an angle $\alpha_{1}$ with $t$, and $v$ is dropped below the plane of the page so that $t_{2}$ makes an angle $\alpha_{2}$ with $t$. We make $\alpha_{2}>\alpha_{1}$ so that the hinge ( $u, e, v$ ) is convex up.

To verify that $t$ can be made locally Delaunay with this configuration, we need to ensure that $S_{\sigma^{\prime \prime}}$ does not contain $u$. Since $\operatorname{aff}(t)=\operatorname{aff}\left(S_{\sigma^{\prime \prime}} \cap S_{\sigma^{\prime}}\right), c_{\sigma^{\prime}}$ and $c_{\sigma^{\prime \prime}}$ will lie on $\ell$, the vertical line through $c_{t}$. Then $t$ will be locally Delaunay iff $c_{\sigma^{\prime}}$ lies above $c_{\sigma^{\prime \prime}}$. This is a standard characterization of the Delaunay property: the circumcentres must be in the same order as the simplices. It can be deduced directly from Lemmas B. 2 and B.3.

We refer to Figure 7.3(b). $c_{\sigma^{\prime}}$ lies on the line $\ell_{1}$ through $c_{t_{1}}$, and perpendicular to $t_{1}$, at the point where this line intersects $\ell$. Likewise, $c_{\sigma^{\prime \prime}}$ lies on $\ell_{2}$, the line perpendicular to $t_{2}$ and through $c_{t_{2}}$. The lines $\ell_{1}$ and $\ell_{2}$ intersect at point $a$, and exploiting the right angles it is easy to show that the angle of intersection is $\alpha_{2}-\alpha_{1}$.

We will have $c_{\sigma^{\prime}}$ above $c_{\sigma^{\prime \prime}}$ if $d_{\mathbb{R}^{3}}\left(c_{e}, a\right)<d_{\mathbb{R}^{3}}\left(c_{e}, c_{\sigma^{\prime}}\right)$. This occurs if

$$
\frac{\delta_{2}}{\sin \left(\alpha_{2}-\alpha_{1}\right)}<\frac{\delta}{\sin \alpha_{1}} .
$$

This condition can easily be obtained if, for example, we make $\delta>\delta_{2}$, and $\alpha_{2}>2 \alpha_{1}$.
Although the above construction demonstrates that the flip-tet of an nlD edge in a cocone mesh need not belong to $\mathcal{D}^{3}(P)$, we have not resolved the question of whether or not a self-Delaunay mesh exists as a cocone mesh. In order for our construction to serve as a five point obstruction to a self-Delaunay mesh existing as a substructure of $\mathcal{D}^{3}(P)$, we would need to verify that the other internal edge $[u, q]$ is locally Delaunay after $e$ has been flipped to $e^{\prime}=[u, v]$. It is not clear that this is possible. In other words, the edge flipping algorithm could go on to flip $[u, q]$ to $[v, w]$ and then flip $e^{\prime}$ to $[p, w]$ so that the final self-Delaunay mesh is a substructure of $\mathcal{D}^{3}(P)$ after all.

### 7.4 Discussion

The observation that a non-sharp mesh cannot have an unflippable nlD edge, complements nicely the fact that a smooth mesh is the desired representation of a smooth surface. It also yields insight into a question recently posed by Agarwal et al. [ASY08]. In that work they considered the problem of "untangling" a planar mesh which results when the vertices of a valid triangulation have been displaced. A mesh is tangled if it contains inverted triangles, where the orientation of the triangles is defined by the counterclockwise vertex ordering in the original triangulation. They mention in their conclusion that it is not clear how to state the tangling problem in the context of surface meshes. At least in the context of meshes representing smooth surfaces, it seems that a sharp mesh is a good candidate for a tangled mesh. Of course, this does not resolve the problem of identifying tangled regions, but it at least allows us to determine whether or not tangling has occured.

Nonetheless, the question of when we can be assured that extrinsic Delaunay edge flipping will not tangle the mesh remains unanswered. We don't know when we can guarantee that a self-Delaunay mesh exists on $P$, let alone whether one exists as a substructure of $\mathcal{D}^{3}(P)$.

## Chapter 8

## Conclusions

Through the exploration of the concept of self-Delaunay meshes we have gained insight into the relationship between various Delaunay structures. In the Euclidean setting, it has been observed that a continuous displacement of the vertices of a Delaunay triangulation results in a change in the graph topology of $\mathcal{D}^{d}(P)$ only if the vertices move through a degenerate configuration [AGMR98]. In other words there is an instant when the vertices are not in general position. Thus a small change in the vertex positions results in a change of the graph topology of $\mathcal{D}^{d}(P)$ only if the initial configuration was close to being not in general position.

The distinction between different Delaunay structures representing a smooth surface, $S$, can similarly be attributed to vertex configurations that are near to being not in general position. In this case it is the metric on $S$ that is undergoing perturbation, rather than the position of the samples. Because such near degenerate configurations cannot be eliminated by any density-based sampling criteria, the iDt-mesh and the rDt are distinct structures and they are not self-Delaunay meshes in general, regardless of sampling density.

In the context of the rDt, configurations of sample points that are near-degenerate are manifest by the presence of sliver tetrahedra in $\mathcal{D}^{3}(P)$. Sliver tetrahedra are a notorious difficulty in Delaunay-based surface reconstruction algorithms, and, without local uniformity constraints on the sampling, Erickson [Eri01] has demonstrated that there may be $\mathcal{O}\left(n^{2}\right)$ sliver tetrahedra in $\mathcal{D}^{3}(P)$. The issue is nicely summarized by Dumitriu et al. [DFKM08], where it is observed that the extraction of a manifold from the collection of cocone triangles is a non-local process. The problem is that there is no canonical global mesh representation of $S$ which does not require explicit knowledge of $S$ itself for its construction. They
demonstrate, however, that on a locally uniform subsample it is possible to locally and conservatively select edges that will belong to a globally consistent polygonal mesh with the topology of $S$. The faces of the mesh are shown to have a vertex size bounded by a constant, and experimentally they rarely have more than five vertices.

Kil and Amenta [KA08] also address the locality problem in their parallelizable surface reconstruction algorithm. They require a local uniformity constraint on the sample set, and like Dumitriu et al., a cornerstone of their algorithm is the identification of edges which unambiguously will belong to the global mesh.

A central problem which has motivated much of the work in this thesis, is the identification of sampling conditions which will guarantee the existence of a smooth self-Delaunay mesh. It is known that a nonobtuse mesh (which is necessarily a self-Delaunay mesh) with vertices on $S$ may be constructed with an arbitrarily fine resolution (i.e., triangle circumradii may be made arbitrarily small). This follows easily from work of Colin de Verdière [CdVM90], for example. However, these existence results only imply that for a given sampling radius there is some point set $P \subset S$, which satisfies the sampling radius and admits an appropriate self-Delaunay mesh ${ }^{1}$. Instead, we are interested in whether there is a sampling radius for which any point set meeting that requirement will admit a smooth self-Delaunay mesh.

The problem remains unresolved, but the core of the issue is similar to the source of the locality problem tackled in [DFKM08] and [KA08], namely near degenerate configurations (i.e., "almost not general position", when there are multiple vertices that are almost cocircular in a given metric), where there is no clear consensus as to what the connectivity should be. I believe that progress in either of these problems will yield progress in the other.

The work of Dumitriu et al. hints that the subgraph, $G$, of the 1-Gabriel complex consisting of those edges which lie close to $S$ and connect geodesically close samples should be sufficiently rich to capture the topology of $S$. In particular the graph would project onto $S$ so as to form the 1-skeleton of a cell complex on $S$ and each 2-cell in $S \backslash G$ would be bounded by a small number of edges projected from $G$. If these observations could be quantified and verified, it could facilitate the task of explicitly constructing a self-Delaunay mesh on $P$. If we start with a smooth mesh homeomorphic to $S$ and containing $G$, then since the Delaunay extrinsic edge flipping algorithm will never flip an edge in $G$, the problem

[^8]of demonstrating that no unflippable edges will be encountered is reduced to demonstrating this only for the regions of disk topology that contain no vertices or edges of $G$ in their interiors.

A deeper understanding of Delaunay extrinsic edge flipping may also yield insight into the locality problem. Since a self-Delaunay mesh on $P$ is not unique, we cannot hope to construct a globally consistent mesh by constructing local self-Delaunay patches. However, it is possible that such an approach would work for meshes with minimal surface area. A mesh with minimal surface area on $P$ will also not be unique in general, but we expect this to be akin to the points not being in general position. In other words, an arbitrarily small perturbation of the vertices would ensure that there is a unique mesh of minimal surface area. It is trivial to extend the Delaunay extrinsic edge flipping algorithm into one which truly minimizes surface area: we simply continue to flip edges as long as the flip will reduce the sum of the subtended angles. If we can show that such edge flipping will not encounter unflippable edges, but instead converges to a smooth mesh of minimal surface area, then such a mesh will be a good candidate for the canonical global mesh surface representation that would present a much simpler resolution to the locality problem than that proposed by Dumitriu et al [DFKM08] or Kil and Amenta [KA08].

It is expected that a mesh with minimal surface area must be smooth if $P$ is well sampled from $S$, but we lack a demonstration. If this could be established, then such a mesh must be (isometric to) a self-Delaunay mesh, since otherwise a non-planar nlD edge could be flipped to reduce the surface area. It is known that the surface area of $S$ may be defined as the infimum of the limit inferior of the sequences of surface areas of all meshes which converge uniformly (i.e. in Hausdorff distance) to $S$ [Tor70, AT72]. However, we have no demonstration that the meshes of minimal surface area on $P$ converge in Hausdorff distance to $S$ as the size of $P$ increases. If this could be established, it would follow immediately that the normal vectors of these meshes must also converge to the nearby normals in $S$ [MT04, HPW06], and hence these meshes would be smooth. These considerations provide reason for optimism that a self-Delaunay mesh that is a good representation of $S$ does exist on $P$, provided $P$ is sufficiently dense.

While it is clear that a smooth mesh that is close in Hausdorff distance to $S$ must also be close in terms of surface normal vectors, it is not clear to what extent the assumption of positional accuracy is redundant. Is a topologically correct $\theta$-smooth mesh, $M$, necessarily a good approximation to $S$ ? By "good approximation" in this context we mean that the
normal error and Hausdorff distance between $M$ and $S$ are $\mathcal{O}(\theta)$. For example, we expect that a smooth and topologically correct Gabriel mesh, if it exists, is a cocone mesh, but we lack a demonstration which permits us to take this for granted. Likewise, the geometric correctness of a smooth self-Delaunay mesh on $P$ requires formal verification. To the best of my knowledge, all mesh structures which currently provide geometric accuracy guarantees with respect to an unknown surface have some constraint on the triangles that is equivalent to a bound on the circumradius. A result quantifying or clarifying the correctness of a smooth mesh would be a powerful generalization.

My interest in self-Delaunay meshes was initially generated by the insights presented by Bobenko and Springborn [BS07] demonstrating the natural affinity between Delaunay triangulations and the cotan operator. However, I now argue that the principle reason that self-Delaunay meshes deserve further study is their promise as a Delaunay-based surface representation that relies upon neither the reference surface, $S$, nor the ambient Delaunay tetrahedralization, $\mathcal{D}^{3}(P)$.

This thesis has presented algorithms for constructing self-Delaunay meshes, but the study of these objects has also yielded many significant incidental contributions. The establishment of the Gabriel complexes as a nested decomposition of the Delaunay triangulation may provide a useful framework for extracting low dimensional pwf manifolds in a high dimensional space. Likewise the intuitive explanation and improvement of the intrinsic sampling criteria developed by Leibon and Letscher [LL00], should serve as a starting point for establishing these criteria in detail for higher dimensional manifolds. The lower bound on the scr in terms of the lfs provides a bridge between the theoretical convenience of intrinsic sampling criteria, and the practical utility of extrinsic sampling criteria.

The relevance of self-Delaunay meshes to geometry processing may depend upon the outcome of further investigations into the Delaunay extrinsic edge flipping algorithm. In the end I feel that the questions raised here and in the end of chapter discussions are amongst the biggest contributions of this thesis.

## Appendix A

## Circumcentric dual cells with negative area

The circumcentric dual complex associated with a triangulation is becoming a structure of interest in discrete differential geometry. It arises naturally in formulations of discrete exterior calculus and, in the two dimensional case that concerns us here, it provides an elegant interpretation of discrete Laplace operators based on the cotangent formula. If the primal triangulation is Delaunay, then the circumcentric dual complex is the Voronoi diagram of the vertices. On the other hand, if a primal edge is not locally Delaunay, the length of the corresponding dual edge will be negative. In many applications this does not present a problem. However in this appendix we draw attention to the possibility that the dual cell to a primal vertex may have negative area, and we discuss some of the implications. We review the definition of circumcentric dual cells and provide simple explicit constructions of triangle configurations in which a primal vertex has a circumcentric dual cell with negative area.

## A. 1 Circumcentric dual cells and their area

The dual structures associated with the simplices in a triangulated manifold play a prominent role in the discrete exterior calculus [Hir03, DHLM05]. In the canonical case where the primal simplices carry a Euclidean metric, the cells that are naturally dual to the primal vertices are the so called circumcentric dual cells. The vertices of the boundaries of these
cells are the circumcentres of the full dimensional simplices that contain the primal vertex. In the cited expositions, the (primal) triangulations under consideration are assumed to be "well-centred", meaning that all simplices contain their circumcentre. This restriction is imposed to simplify the exposition, and indeed in this case, the dual cells of the primal vertices are simply their Voronoi cells. In fact this is true if the triangulation is Delaunay, which is a much weaker constraint than the well-centred criterion.

We will focus our attention on the two dimensional case. (A manifold triangle mesh is a two dimensional object. The fact that it is immersed in three dimensions does not make it a three dimensional mesh!) So demanding a well-centred triangulation is equivalent to demanding a nonobtuse triangulation. For a fixed set of vertices (cone points), such a triangulation of a piecewise flat surface will not exist in general. For example, a planar quadrilateral with two adjacent obtuse interior angles admits no nonobtuse triangulation on its four vertices.

Anyway, we are here interested in the case of arbitrary triangulations. In this case the dual cells become more complicated. They can "fold over" on themselves. The circumcentric dual cells are a special case of orthogonal duality structures that have been studied in detail by Glickenstein [Gli05].

For simplicity, we will work with a planar triangulation. Let the primal vertex under consideration be $p$. The triangles which contain $p$ form the star of $p$. The dual cell to $p$, denoted $\star p$ is a polygon whose vertices are the circumcentres of the triangles in the star of $p$. If $t_{1}$ and $t_{2}$ are adjacent triangles that share edge $e$, then the line segment connecting the circumcentres of $t_{1}$ and $t_{2}$ (denoted $c_{1}$ and $c_{2}$ ) will lie on the perpendicular bisector of $e$. This line segment, denoted $\star e$, is dual to $e$ and belongs to the formal boundary of $\star p$. See Figure A.1.

We will describe precisely how to compute the area of a circumcentric dual cell below, but informally, we can compute the area by cutting the cell into triangular pie slices such that the base of a given slice is a dual edge, and the point of the slice is the primal vertex $p$. The area of the cell is then given by the sum of the areas of the triangular pie slices. The twist is that some of the dual edges can have negative length, and in this case the corresponding pie slice has negative area. The negative length dual edges occur when the rotational ordering around $p$ of the corresponding circumcentres is reversed with respect to that of their primal triangles. This occurs precisely when the primal edge is not locally Delaunay. So for example, in Figure A.1, the primal edge $[p, v]$ is not locally Delaunay,


Figure A.1: A dual cell folds on itself when one of the primal edges incident to the vertex $p$ is not locally Delaunay. In this case $[p, v]$ is nlD and so its dual edge, $\left[c_{1}, c_{2}\right]$, has a negative length.
so its dual edge $\left[c_{2}, c_{1}\right]$ has negative length. The area of the dual cell is the sum of the areas of seven triangular pie slices, but the slice $\left[p, c_{1}, c_{2}\right]$ has a negative area. Portions of contributions of this slice are cancelled out by the positive area contributions of slices [ $\left.p, c_{0}, c_{1}\right]$ and $\left[p, c_{2}, c_{3}\right]$. The net effect is that the dual cell has a region of positive area around $p$ and a small triangle of negative area below $\left[c_{2}, c_{1}\right]$.

The definition of the signed length of a dual edge, the area of the whole dual cell, and how to compute them is as follows (See [Gli05][p.10]). Let $e=[p, v]$ be an edge with adjacent triangles $t_{1}$ and $t_{2}$. Let $c_{1}$ and $c_{2}$ be the circumcentres of $t_{1}$ and $t_{2}$ respectively, and let $c_{e}$ be the midpoint of $e$. Then the signed length of $\star e$ is given by

$$
|\star e|=d_{ \pm}\left(c_{1}, c_{e}\right)+d_{ \pm}\left(c_{2}, c_{e}\right),
$$

where $d_{ \pm}\left(c_{i}, c_{e}\right)$ is equal to the Euclidean distance between $c_{i}$ and $c_{e}$ with sign positive if $c_{i}$ is on the same side of the line supporting $e$ as is $t_{i}$, and sign negative if $c_{i}$ is on the opposite side.

The area contribution to $\star p$ from $\star e$ is given by the sum of the signed areas of the two right triangles $\left[c_{e}, p, c_{1}\right]$ and $\left[c_{e}, p, c_{2}\right]$, where the sign of $\left|\left[c_{e}, p, c_{i}\right]\right|$ is the same as the sign of


Figure A.2: Edge $[p, v]$ is locally Delaunay (as evidenced by the colourful circumcircles), so $\left[c_{1}, c_{2}\right]=\star[p, v]$ has positive length. The contribution to the area of $\star p$ from $\left[c_{1}, c_{2}\right]$ is given by the sum of the signed areas of triangles $\left[c_{e}, p, c_{1}\right]$ and $\left[c_{e}, p, c_{2}\right]$. The former has a negative area, but the latter has a positive area of larger magnitude.
$d_{ \pm}\left(c_{i}, c_{e}\right)$. See Figure A. 2 for an example. The area of $\star p$ is the sum of the area contributions of all the dual edges.

## A. 2 Construction of a dual cell with negative area

Upon examining Figure A. 1 a method of constructing the star of $p$ such that $\star p$ has negative area presents itself: If we hold $p$ and all other vertices except for $v$ fixed, and move $v$ such that edge $[p, v]$ grows in length but maintains its orientation, then $\left[c_{1}, c_{2}\right]$ will move away from $p$, since it must always lie on the perpendicular bisector of $[p, v]$. Also since $c_{1}$ and $c_{2}$ must lie on the perpendicular bisectors of $\left[p, v_{1}\right]$ and $\left[p, v_{2}\right]$ respectively, and these bisector lines remain fixed, $\left[c_{1}, c_{2}\right]$ will grow in length. Thus the triangle of negative area with base $\left[c_{1}, c_{2}\right]$ will grow as $v$ moves away from $p$, and eventually this negative area will dominate the area of $\star p$.

Thus we can create a cell with negative area by introducing a single not locally Delaunay edge, but the ratio of the longest to shortest edges incident to $p$ becomes large in this case. We can create a less pathological star with an associated dual cell of negative area if we


Figure A.3: (a) Triangle $t=\left[p, v, v_{1}\right]$ is drawn in its circumcircle. The contribution of $t$ to the area of $\star p$ is the negative area triangle $\left[p, c_{e_{1}}, c_{t}\right]$ and the positive area triangle $\left[p, c_{e}, c_{t}\right]$. These triangles share a common base $\left[p, c_{t}\right]$ which is also the diameter of their common circumcircle. If we arrange $\left[p, v_{1}\right]$ to make an angle of $\pi / 4$ with $\left[p, c_{t}\right]$, then $c_{e_{1}}$ has a maximal distance from $\left[p, c_{t}\right]$ and so if $[p, v]$ makes a larger angle with $\left[p, c_{t}\right], t$ will necessarily have a net negative area contribution to the dual cell. (b) We construct such a $t$ to have an angle of $\pi / 5$ at $p$ and then construct the star of $p$ out of copies of this triangle and its mirror as shown (sketched - the actual object is probably pointier). The dual cell of such a star will have a negative area.
include several not locally Delaunay edges incident to the central vertex.
The computation of the area contribution of a dual edge is broken down into two right triangular contributions, one for each of the two circumcentres associated to triangles incident to the primal edge. This breakdown allows us to consider the contribution from each primal triangle to the area of the dual cell, rather than using edges as primitives. We will construct a triangle that has a net negative contribution to the dual cell area and build the star of $p$ out of such triangles.

Referring to Figure A.3(a), the primal triangle $t=\left[p, v, v_{1}\right]$ contributes two right triangles, $\left[p, c_{e}, c_{t}\right]$ and $\left[p, c_{e_{1}}, c_{t}\right]$, to the area of the dual cell, $\star p$. The common base, $\left[p, c_{t}\right]$, of these triangles is the diameter of their common circumcircle. It is easy to see that the primal edges $[p, v]$ and $\left[p, v_{1}\right]$ can be manipulated so as to make the negative area triangle $\left[p, c_{e_{1}}, c_{t}\right]$ larger than the positive area triangle $\left[p, c_{e}, c_{t}\right]$. There is quite a bit of freedom and
we can constrain the angle at $p$ of such a primal triangle to be $\pi / 5$. In this way we can construct a star around $p$ using five copies of this triangle and five copies of its mirror, as shown in Figure A.3(b). The resulting star will have a negative area dual cell.

## A. 3 Implications for the DEC on arbitrary meshes

The orthogonal duality structures provide an elegant framework which allowed Glickenstein to define Laplacian operators on arbitrary triangulations. In [Gli05][eq (15)], the Laplacian operator suggested by [Hir03] is written as

$$
\begin{equation*}
(\Delta f)_{i}=\frac{1}{|\star\{i\}|} \sum_{j \in N(i)} \frac{|\star\{i, j\}|}{|\{i, j\}|}\left(f_{j}-f_{i}\right), \tag{A.1}
\end{equation*}
$$

where $N(i)$ is the set of vertices in the triangulation that share an edge with $\{i\}$. Glickenstein then suggests the following alternative definition:

$$
\begin{equation*}
(\Delta f)_{i}|\star\{i\}|=\sum_{j \in N(i)} \frac{|\star\{i, j\}|}{|\{i, j\}|}\left(f_{j}-f_{i}\right) . \tag{A.2}
\end{equation*}
$$

The motivation for the subtle difference is clear. Where Hirani restricted the exposition to well centred triangulations, Glickenstein is allowing for more general triangulations. We have seen that in this case $|\star\{i\}|$ may be negative. It follows easily that it may also vanish: a continuous deformation argument applies to our construction from Figure A.1, or the construction of Figure A. 3 can be modified to use primal triangles that have a net contribution of zero to the dual cell area.

However, for many applications (e.g., spectral processing [VL08]) we desire a Laplacian that, when operating on discrete functions defined on the primal vertices, yields exactly such functions. The DEC is designed to provide such Laplacian operators. The distinction between Equation (A.1) and Equation (A.2) is more than a mere formality. For example, in spectral processing, the difficulties cannot be escaped by employing a generalized eigenvalue problem, as is normally done in the FEM. The problem lies in the associated mass matrix not being positive definite. In other words the difficulty is not limited to dual cells with vanishing area. If we are to employ a Laplacian in the form of Equation (A.1), we need to define an inner product on our discrete function space using a quadrature that renders our Laplacian self-adjoint (this is essentially the "discrete Hodge star" for functions). This is
attained when we define

$$
\begin{equation*}
\langle f, g\rangle=\sum_{i \in V} f_{i} g_{i}|\star\{i\}| . \tag{A.3}
\end{equation*}
$$

So our inner product is defined by a bilinear form represented by a diagonal matrix whose entries are the areas of the corresponding dual cells. If we have negative cell areas this matrix is not positive definite. It is not clear how to view this situation as anything other than undesirable.

For a general triangulation, an orthogonal duality structure may not be the appropriate choice. As pointed out in [MDSB03], the cotan weights can be derived from a diffusion formulation assuming any dual cell whose boundary intersects the primal edges at their midpoints. An appropriate choice may be the barycentric cells. In this case our quadrature (A.3) amounts to the lumped mass approximation in the FEM.

The use of the Voronoi cells as suggested in [MDSB03] is not appealing for general meshes since the neighbour relations of these cells do not reflect the connectivity of the triangulation (and hence of the Laplacian). Their argument for the tight error bounds resulting from the use of Voronoi cells is not compelling because piecewise constant functions are assumed (constant function value on each Voronoi cell), whereas piecewise linear functions are employed for all the derivations.

Orthogonal duality structures provide a particularly elegant formulation of the DEC. However it appears that there is a limited class of triangulations for which such structures may be reliably employed. In a footnote, Elcott and Schröder [ES06] go so far as to say that circumcentric dual cells are only appropriate when a Delaunay triangulation is used. Certainly, we need a Delaunay triangulation if we want a Free Lunch [WMKG07], but in general, the restriction to Delaunay triangulations, at least in the two dimensional case, is somewhat stronger than is necessary for many applications. We have the following:

1. For positive length circumcentric dual edges, a Delaunay triangulation is necessary and sufficient.
2. For positive area circumcentric dual cells, a Delaunay triangulation is sufficient.

In terms of the Laplacian operator, positive length circumcentric dual edges equates to positive weights in the cotan operator. However, when the quadrature (A.3) is employed, the operator (A.1) remains positive semi-definite on arbitrary triangulations (shown by means of the discrete Dirichlet energy [PP93]). We don't always need to demand positive
weights. However, if our dual cells admit negative areas, the situation really does start to look unpleasant.

The area of a dual cell can be expressed in terms of the primal triangulation by means of the cotan formula. Letting $E(i)$ denote the set of edges incident to $\{i\}$, and expressing the cotan formula as

$$
\frac{|\star e|}{|e|}=\frac{1}{2}\left(\cot \alpha_{e}+\cot \beta_{e}\right),
$$

we have

$$
\begin{equation*}
|\star\{i\}|=\frac{1}{4} \sum_{e \in E(i)}|\star e||e|=\frac{1}{8} \sum_{e \in E(i)}\left(\cot \alpha_{e}+\cot \beta_{e}\right)|e|^{2} . \tag{A.4}
\end{equation*}
$$

The sum of the areas of the dual cells is equal to the area of the domain [Gli05][Proposition 9]. In the case where our triangulation describes a mesh, i.e., each primal simplex is a Euclidean simplex in the ambient embedding space, then Equation (A.4) reveals that $\sum_{i}|\star\{i\}|$ is the discrete Dirichlet energy of the embedding map [PP93].

Demanding that the circumcentric dual cells have positive area is weaker than demanding that the triangulation be Delaunay. We are demanding that in some sense the triangulation is locally never too far from being Delaunay.

## Appendix B

## Sphere lemmas

For convenience we provide here some standard results on spheres. In Section B. 1 we provide a series of similar lemmas regarding the intersections of spheres of an arbitrary common dimension. In Section B. 2 we provide a couple of observations on small spherical triangles.

## B. 1 Sphere intersections

Here we consider the intersection of $S_{1}$ and $S_{2}$, two $k$-spheres in $\mathbb{R}^{k+1}, k \geq 1$, with centres and radii $c_{1}$ and $r_{1}$, and $c_{2}$ and $r_{2}$ respectively. Throughout this section, $\ell$, denotes the line containing $c_{1}$ and $c_{2}$.

We present a number of similar lemmas which highlight how specific properties of the intersection relate to the relative sizes of the spheres. The principle observation is that if the claim is true for $k=1$, then it will be true for all $k \geq 1$, because any planar slice that includes the sphere centres will look the same, as depicted in Figure B.1, for example.

Lemma B. 1 If two $k$-spheres intersect tangentially, their intersection is a single point. If they intersect transversely, their intersection is a $k-1$-sphere whose affine hull is orthogonal to the line containing the centres of the spheres.

Proof Suppose $S_{1}$ and $S_{2}$ intersect tangentially, and $x \in S_{1} \cap S_{2}$. Then $T_{x} S_{1}=T_{x} S_{2}$ and since $\overrightarrow{c_{1} x}$ and $\overrightarrow{c_{2} x}$ are perpendicular to the tangent space at $x, x$ must lie on $\ell$. Since there is only one point on $\ell$ whose distance from $c_{1}$ is $r_{1}$ and whose distance from $c_{2}$ is $r_{2}$, it follows that $x$ is the unique point in $S_{1} \cap S_{2}$.


Figure B.1: The intersection between two $k$-spheres is characterized by the intersection of two circles. For any point $x \in \operatorname{aff}\left(S_{1} \cap S_{2}\right)$, consider the circles defined by the intersection of $S_{1}$ and $S_{2}$ with the plane defined by $c_{1}, c_{2}$, and $x$. For all choices of $x$ we get the same picture; the same triangle $\left[x, c_{1}, c_{2}\right]$.

Now suppose that $S_{1}$ and $S_{2}$ intersect transversely, and again let $x \in S_{1} \cap S_{2}$. Then consider the triangle $\left[x, c_{1}, c_{2}\right]$, and let $o$ be the foot of its altitude on base $\left[c_{1}, c_{2}\right]$, as shown in Figure B.1. Now, for any other point $y \in S_{1} \cap S_{2}$, the triangle $y, c_{1}, c_{2}$ ] will be congruent to $\left[x, c_{1}, c_{2}\right]$, since corresponding sides have equal length. Thus $o$ is the common foot of all such triangles on base $\left[c_{1}, c_{2}\right.$ ], and the distance from $o$ to any $x \in S_{1} \cap S_{2}$ is constant, equal to the height of the triangle. Also, by construction of $o$, the vector $\overrightarrow{o x}$ is perpendicular to $\ell$, for any $x \in S_{1} \cap S_{2}$. It follows that $S_{1} \cap S_{2}$ is a sphere with centre $o$ and whose affine hull is orthogonal to $\ell$.

Lemma B. 2 If two spheres intersect then the circumcentre of the larger sphere is farther from the affine hull of their intersection.

Proof Since $d_{\mathbb{R}^{k+1}}\left(c_{i}, \operatorname{aff}\left(S_{1} \cap S_{2}\right)\right)=d_{\mathbb{R}^{k+1}}\left(c_{i}, o\right)$, the lemma follows from the Pythagorean Theorem: $r_{i}^{2}=d_{\mathbb{R}^{k+1}}\left(c_{i}, o\right)^{2}+|[o, x]|^{2}$, where $i=1,2$ and $x \in S_{1} \cap S_{2}$, as in Figure B.1.

Lemma B. 3 If $S_{1}$ and $S_{2}$, intersect transversely, then on one side of the affine hull of their intersection, $S_{1}$ contains $S_{2}$, and on the other side $S_{2}$ contains $S_{1}$. Further, if the centres of both spheres lie on the same side, then the sphere with the larger radius contains the other on that side.

Proof The first statement follows from Lemma B.1, since if $x, y \in S_{1}$ are on the same side of aff ( $S_{1} \cap S_{2}$ ), then they can be connected by an curve on $S_{1}$ which does not pass through $\operatorname{aff}\left(S_{1} \cap S_{2}\right)$, and therefore does not intersect $S_{2}$.

For the second statement, suppose that $r_{1}>r_{2}$. Then it follows that $c_{1}$ is farther from aff $S_{1} \cap S_{2}$ than is $c_{2}$. Then by considering the intersection of $\ell$ with the spheres, it is immediate that $S_{1}$ contains $S_{2}$ on the side of aff $\left(S_{1} \cap S_{2}\right)$ that contains the centres.

Lemma B. 4 If $S_{1}$ and $S_{2}$ intersect transversely such that the centre of $S_{2}$ lies in the affine hull of their intersection, then $S_{2}$ has a smaller radius than $S_{1}$ and it is contained in $S_{1}$ on the side of aff $\left(S_{1} \cap S_{2}\right)$ that contains the centre of $S_{1}$.

Proof The proof is essentially the same as that of Lemma B.3.
Lemma B. 5 If $S_{1}$ and $S_{2}$ intersect transversely and if $S_{1}$ contains $S_{2}$ on the side of aff ( $S_{1} \cap$ $S_{2}$ ) that contains $c_{2}$, then $S_{1}$ is larger than $S_{2}$.

Proof Let $z$ be the point where $\ell$ intersects $S_{2}$ on the side of aff $\left(S_{1} \cap S_{2}\right)$ containing $c_{2}$. Suppose that $c_{1}$ were on the opposite side of aff $\left(S_{1} \cap S_{2}\right)$. Then $\left[c_{2}, z\right]$ would be contained in a radius of $S_{1}$, and therefore $r_{1}>r_{2}$. However, if this were the case, then $S_{1}$ would also contain $S_{2}$ on the side of aff $\left(S_{1} \cap S_{2}\right)$ containing $c_{1}$, since $\ell$ would have to intersect $S_{1}$ farther from $\operatorname{aff}\left(S_{1} \cap S_{2}\right)$ on that side. This violates the first statement of Lemma B.3, and the hypothesis that the spheres intersect transversely.

Therefore, $c_{1}$ must lie on the same side of $\operatorname{aff}\left(S_{1} \cap S_{2}\right)$ as $c_{2}$, and the lemma follows from the second statement of Lemma B.3.

## B. 2 Spherical geometry

We present here a couple of observations about spherical triangles. By means of the Gauss map these results translate directly into observations about valence three vertices in a mesh. A spherical triangle is a region on a unit hemisphere bounded by three minimal geodesic arcs
connecting three distinct points. Any three distinct points on the sphere that are pairwise not antipodal and which do not lie on a common great circle, can be contained in an open hemisphere. These points define a unique spherical triangle, the edges of which all have length less than $\pi$.

We make use of two basic tools of spherical geometry. A lune is a portion of the sphere contained between two geodesic arcs meeting at antipodal points, as depicted in Figure B.2. Thus a lune is a spherical polygon with two sides. If the angle between the geodesics at a vertex is $\eta$, then, since the area of the sphere is $4 \pi$, the area of the lune is $\frac{\eta}{2 \pi} 4 \pi=2 \eta$.

The other observation we will need is Girard's Theorem, which states that the area of a triangle $t$ with angles $\eta, \theta$, and $\varphi$ is given by

$$
|t|=\eta+\theta+\varphi-\pi .
$$

Girard's theorem can be proven by considering the lunes defined by the sides of $t$, as well as their counterparts formed by the other half of the great circles.


Figure B.2: A spherical triangle in the lune defined by two of its sides.

Lemma B. 6 If every side in a spherical triangle has length less than $\pi / 2$, then the sum of any two angles of the triangle is less than $\pi$.

Proof Let $\eta, \theta$, and $\varphi$ be the angles of the spherical triangle $t$ whose sides are each less than $\pi / 2$. Consider the lune defined by the great circles that support the sides that define the angle $\eta$ (Figure B.2). The area of this lune is $2 \eta$. Since the sides of $t$ have length less than $\pi / 2, t$ is contained in one half of the lune. Therefore, by Girard's theorem for the area of a spherical triangle we have

$$
|t|=\eta+\theta+\varphi-\pi<\eta .
$$

Thus $\theta+\varphi<\pi$. The other angle sums are demonstrated in the same way using the appropriate choice of lune.

The same argument applies if we allow the side lengths to be equal to $\pi / 2$ :
Lemma B. 7 If every side in a spherical triangle has length less than or equal to $\pi / 2$, then the sum of any two angles of the triangle is less than or equal to $\pi$.

An equivalent statement to Lemma B. 6 is that if a valence three vertex, $p$, in a triangle mesh is non-sharp, then the sum of any two of of the face angles incident to $p$ is more than $\pi$. The Gauss map can be used to establish the equivalence of the statements. This means that the hinge of any edge incident to $p$ unfolds to a non-convex quadrilateral. In particular, there are no nlD edges incident to $p$.

The converse to Lemma B. 6 does not hold. However, we do have Lemma B.8, below, which says that if all the face angles incident to a valence three vertex are obtuse, then the vertex is non-sharp.

Lemma B. 8 In a spherical triangle, if every angle is acute, then every side has length less than $\pi / 2$.

Proof Let $t=\{a, b, c\}$ be a spherical triangle with angles $\eta, \theta$ and $\varphi$. First we observe that $t$ cannot have two sides with lengths $\geq \frac{\pi}{2}$. Suppose to the contrary that $|a b| \geq \frac{\pi}{2}$ and $|a c| \geq \frac{\pi}{2}$. Then consider the lune with angle $\eta$ (at $a$ ). Since the area of the lune is $2 \eta$, and $t$ fills more than half the lune, Girard's theorem yields

$$
|t|=\eta+\theta+\varphi-\pi \geq \eta
$$

But that implies that $\theta+\varphi \geq \pi$, contradicting the hypothesis that the angles are acute.
Therefore at least two of the sides have length less than $\frac{\pi}{2}$. Again assume that these sides meet at $a$, and consider the lune with angle $\eta$. Since $t$ occupies less than half the lune, and $\eta<\frac{\pi}{2}$, it follows that $t$ is contained in an octant of the sphere.

Any geodesic contained in the interior of an octant will have length less than $\pi / 2$. To see this, consider the octant $\Omega$ defined by the positive coordinate octant in $\mathbb{R}^{3}$. Any two points $p, q \in \Omega$ may be considered as unit vectors, and their scalar product will be positive. This implies that the angle between them is less than $\frac{\pi}{2}$, but the angle between them is precisely the length of the geodesic connecting them in $\Omega$.

Therefore, all sides of $t$ must have length less than $\frac{\pi}{2}$.

## Appendix C

## More on the intrinsic Voronoi diagram

This appendix presents material to supplement Chapter 5.
In Section C. 1 we discuss the assumption that the sample points are in general position with respect to the intrinsic metric of the surface. The conclusion is that the assumption is reasonable in the context presented in Chapter 5, but it should not have been taken for granted that this must be so. In particular, this issue will need to be addressed with care if an attempt is made to extend the results of Chapter 5 to higher dimensional manifolds.

Section C. 2 gives the background and intuition behind Lemma 5.18. Although he did not explicitly state the lemma, its proof is due to Chavel [Cha06], and we provide a sketch of this proof.

In Section C. 3 we present a geometrically intuitive conjecture that would enable us to significantly improve the bound of Theorem 5.19 and Corollary 5.20.

## C. 1 On the assumption of general position

In Chapter 5, we made the assumption that the sample set $P \subset S$ is in general position with respect to the intrinsic metric $d_{S}$ on $S$. In this section we show that the issue is delicate and deserves more attention than we gave it. We give an argument to support our assumption of general position in the context of Chapter 5, but the argument is not rigorous, and it does not apply to higher dimensional manifolds or surfaces that are not compact.

The assumption of general position is natural and well accepted when doing theoretical work with Delaunay complexes in Euclidean space. In practice, general position can be simulated symbolically [EM90]; the infinitesimal perturbation of the points that may be needed to achieve general position does not need to be performed explicitly.

In their development of the closed ball property criterion, Edelsbrunner and Shah [ES94] worked with the rVd , and in that context the assumption of general position with respect to $\left.d_{\mathbb{R}^{3}}\right|_{S \times S}$ was expressed as a general intersection property, which demands that Voronoi facets of the ambient Euclidean Voronoi diagram, if they intersect $S$, do so in an appropriate way. Specifically, Voronoi faces intersect $S$ in one dimensional submanifolds (a collection of Voronoi edges in the rVd), Voronoi edges intersect $S$ at points (Voronoi vertices in the rVd ), and Voronoi vertices do not intersect $S$. Morse-theoretic arguments were given to show that this assumption is reasonable. An arbitrarily small perturbation will ensure that the points have the general intersection property with respect to $S$.

In our case we are using $d_{S}$, the intrinsic metric on $S$, and the general intersection property does not apply. Instead, we adopt the viewpoint that the samples are in general position if every Voronoi vertex has exactly three incident Voronoi edges. In other words, every Voronoi vertex is associated with exactly three distinct sample points. We call these atomic Voronoi vertices. If a Voronoi vertex, $v$, is associated with $k>3$ (not necessarily distinct) sample points, then an infinitesimal perturbation of those points will suffice to split $v$ into $k-2$ separate atomic Voronoi vertices. We may say that $v$ consists of $k-2$ atomic Voronoi vertices.

From this perspective, the demand that $P$ be in general position is reasonable provided that there are only a finite number of atomic Voronoi vertices. We argue that if there are an infinite number of atomic Voronoi vertices, then $S$ must have infinite genus, and therefore cannot be compact ${ }^{1}$.

If there are an infinite number of atomic Voronoi vertices, then there must be two Voronoi cells, $\mathcal{V}(p)$, and $\mathcal{V}(q)$, that share an infinite number of atomic Voronoi vertices. The situation is represented schematically in Figure C.1, where an infinite sequence of Voronoi edges between $\mathcal{V}(p)$ and $\mathcal{V}(q)$ delimits an infinite number of regions, $\Omega_{k}$, composed of Voronoi cells distinct from $\mathcal{V}(p)$ and $\mathcal{V}(q)$. (To accommodate atomic Voronoi vertices that coincide, i.e., that belong to a single non-atomic Voronoi vertex, a Voronoi edge connecting $\Omega_{k}$ with

[^9]

Figure C.1: The boundary between Voronoi cells $\mathcal{V}(p)$ and $\mathcal{V}(q)$ which share an infinite number of Voronoi vertices is represented schematically by a "pearl necklace" of regions $\Omega_{k}$ composed of Voronoi cells distinct from $\mathcal{V}(p)$ and $\mathcal{V}(q)$, strung out between the Voronoi edges separating $\mathcal{V}(p)$ and $\mathcal{V}(q)$. Each region $\Omega_{k}$ contributes two Voronoi vertices that are shared by $\mathcal{V}(p)$ and $\mathcal{V}(q)$, but as described in the text, $\Omega_{k}$ and $\Omega_{j}$ may just represent different boundary components of a common component of $S \backslash(\mathcal{V}(p) \cup \mathcal{V}(q))$.
$\Omega_{k+1}$ may have length zero.) Now since $|P|$ is finite, the $\Omega_{k}$ cannot all be distinct. In fact, there must be a finite number of such components in $S \backslash(\mathcal{V}(p) \cup \mathcal{V}(q))$.

Since an infinite number of the $\Omega_{k}$ must be identified, it follows that $S$ must have infinite genus. Indeed, one of the components of $S \backslash(\mathcal{V}(p) \cup \mathcal{V}(q))$, call it $\Omega_{*}$, has an infinite number of boundary components, each homeomorphic to $S^{1}$, the circle. Thus $S$ could be constructed by performing surgery on a sphere, and part of this surgery would involve cutting an infinite number of disks out of the sphere, and to the boundary of these holes attaching the boundary components of $\Omega_{*}$. This creates an infinite number of handles on $S$.

Fortunately, the assumption that $S$ is compact has eliminated such objects. A surface with infinite genus cannot be compact. This can be seen by observing for example that if $S$ is contained in a bounded volume, then the handles of $S$ must accumulate somewhere, in the sense that there must be a point $x \in \mathbb{R}^{3} \backslash S$ which has no open neighbourhood that does not intersect $S$. Alternatively, we could argue that there must be arbitrarily thin handles, and so the mean curvature is not bounded on $S$, violating the extreme value theorem.

In order to develop the intuition presented here into a more rigorous argument for the validity of the assumption of general position, we would require a more careful definition of atomic vertices. Then we would seek to show that the number of atomic vertices that could be shared by any given three samples is finite, and depends on the genus of $S$. However, the validity of the assumption of general position in higher dimensional manifolds is also a question of interest, and the above topological argument is unlikely to apply in this case. Instead an argument based on bounded curvature would be more likely to succeed.

## C. 2 Chavel's convexity lemma

The proof of the bound on the scr with respect to the lfs depends crucially on the Chavel Lemma 5.18, which we restate here for convenience:

Lemma C. 1 (Chavel) Let $x p q$ be a geodesic triangle consisting of geodesics $\gamma_{1}, \gamma_{2}$ and $\gamma_{3}$, connecting $q$ with $x, x$ with $p$, and $p$ with $q$ respectively. Suppose there are constants $\rho_{0}, r_{0}>0$ with $\rho_{0} \leq \rho_{i}(x)$, and such that $\rho_{G}(z) \geq r_{0}$ on $B_{S}\left(x ; \rho_{0}\right)$. If

$$
\sum_{i=1}^{3} \ell\left(\gamma_{i}\right)<\min \left\{2 \rho_{0}, 2 \pi r_{0}\right\}
$$

then

$$
\gamma_{3} \subset B_{S}(x ; \rho)
$$

where $\rho=\max \left\{d_{S}(x, p), d_{S}(x, q)\right\}$.
The lemma was never explicitly stated by Chavel, but rather it is implicit in the proof of a theorem, [Cha06][IX.6.1, p.404], which gives a global lower bound for the scr in terms of global bounds on the injectivity radius and the Gaussian curvature. Our Theorem 5.19 is essentially a local version of Chavel's theorem and its proof does not employ new insights beyond those that led to the proof of Theorem 5.17. Our purpose here is to describe the main ideas behind Lemma C.1, and to provide an informal outline of its proof as described in [Cha06].

## C.2.1 Overview

We are interested in the strong convexity radius at a point $x$, and we use Chavel's lemma to ensure axiom 2; that the minimal geodesic between two points in the disk lies entirely
within the disk. We consider a geodesic triangle $x p q$ and the lemma demands two bounds on the length of the perimeter of $x p q$. First, it must be less than 2
$r h o_{0}<2 \rho_{i}(x)$. This bound ensures that all points on the triangle are contained within the injectivity radius of $x$. This ensures that the triangle can be lifted onto $T_{x} S$ via the exponential map, which is diffeomorphic on $B_{T_{x} S}\left(0 ; \rho_{i}(x)\right)$. Also, the lengths of $x p$ and $x q$ are preserved under this operation.

The demonstration of the lemma makes use of the Rauch comparison theorem, which enables a comparison between geodesics on different surfaces when the curvature of one is bounded with respect to the curvature of the other. The idea here is that we lift the geodesic triangle $x p q$ onto $T_{x} S$, and then identify $T_{x} S$ with $T_{\tilde{x}} \tilde{S}$, where $\tilde{S}$ is a sphere of radius $r_{0}$, and $\tilde{x}$ is a point on $\tilde{S}$. Because of the curvature bound $\rho_{G}(y) \geq r_{0}$ on $B_{S}\left(x ; \rho_{0}\right)$, we are assured that $\tilde{S}$ has a greater curvature than $S$ for all points on the triangle. The second bound on the perimeter of $x p q$, that it be less than $2 \pi r_{0}$, ensures that the lifted triangle lies within the injectivity radius of $\tilde{x}$. This ensures that the exponential map onto $\tilde{S}$ is diffeomorphic on a neighbourhood containing the lifted triangle, which is required for the employment of the Rauch theorem.

The Rauch theorem enables us to compare the shape of the triangle $x p q$ with that of a corresponding triangle $\tilde{x} \tilde{p} \tilde{q}$ on $\tilde{S}$. Speaking very roughly, Chavel's lemma is saying that the sides of $x p q$ don't bulge out more than the sides of $\tilde{x} \tilde{p} \tilde{q}$. Specifically, if the conditions of the lemma are satisfied, then a minimal geodesic between $p$ and $q$ will be contained within any disk centred at $x$ and containing both $p$ and $q$, which is what we need.

## C.2.2 An outline of the proof

The proof relies on an exercise, [Cha06] [IX.1, p.420], which we state here in our own notation. A geodesic hinge at $x \in S$, denoted $\left(\gamma_{1}, \gamma_{2}, \alpha\right)$ is a curve consisting of two geodesic segments, $\gamma_{1}, \gamma_{2}$, that meet at $x$ with an angle $\alpha$ (see Figure C.2(a)).

Lemma C. 2 (Chavel, Exercise IX.1) Let $\left(\gamma_{1}, \gamma_{2}, \alpha\right)$ be a geodesic hinge at $x \in S$ and


Figure C.2: The geodesic hinge shown in (a), together with $\gamma_{3}$, a minimal geodesic from $p=\gamma_{2}\left(\ell_{2}\right)$ to $q=\gamma_{1}(0)$, forms the geodesic triangle drawn in black in (b). Roughly speaking, Rauch's theorem asserts that geodesics spread apart faster on a manifold with a smaller curvature. Thus, provided there are no conjugate points on the long sides, the short side on the infinitesimal triangle depicted in red will be longer in the manifold with smaller curvature. Lemma C. 3 is an integrated version of the Rauch theorem.
let $r_{0} \geq 0$ be a constant such that

$$
\begin{gathered}
\rho_{G}(y) \geq r_{0} \quad \forall y \in B_{S}\left(x ; \rho_{i}(x)\right) \\
\quad \text { and } \\
\sum_{i=0}^{3} \ell_{i}<\min \left\{2 \rho_{i}(x), 2 \pi r_{0}\right\},
\end{gathered}
$$

where $\ell_{i}=\ell\left(\gamma_{i}\right)$, for $i=1,2$ and $\ell_{3}$ is the geodesic distance between the endpoints of the hinge $\left(\gamma_{1}(0)\right.$ and $\left.\gamma_{2}\left(\ell_{2}\right)\right)$.

Let $\tilde{S}$ be the sphere of radius $r_{0}$, and let $\left(\tilde{\gamma}_{1}, \tilde{\gamma}_{2}, \alpha\right)$ be a geodesic hinge in $\tilde{S}$ with $\ell\left(\tilde{\gamma}_{i}\right)=\ell_{i}$. Then

$$
d_{S}\left(\gamma_{1}(0), \gamma_{2}\left(\ell_{2}\right)\right) \geq d_{\tilde{S}}\left(\tilde{\gamma}_{1}(0), \tilde{\gamma}_{2}\left(\ell_{2}\right)\right) .
$$

For the proof of Lemma C. 1 we will use a slightly more local statement of this exercise:
Lemma C. 3 (Chavel, Exercise IX.1, localized) Let ( $\gamma_{1}, \gamma_{2}, \alpha$ ) be a geodesic hinge at
$x \in S$ and let $r_{0}, \rho_{0} \geq 0$, be constants such that $\rho_{0} \leq \rho_{i}(x)$,

$$
\begin{gathered}
\rho_{G}(y) \geq r_{0} \quad \forall y \in B_{S}\left(x ; \rho_{0}\right) \\
\quad \text { and } \\
\sum_{i=0}^{3} \ell_{i}<\min \left\{2 \rho_{0}, 2 \pi r_{0}\right\},
\end{gathered}
$$

where $\ell_{i}=\ell\left(\gamma_{i}\right)$, for $i=1,2$ and $\ell_{3}$ is the geodesic distance between the endpoints of the hinge $\left(\gamma_{1}(0)\right.$ and $\left.\gamma_{2}\left(\ell_{2}\right)\right)$.

Let $\tilde{S}$ be the sphere of radius $r_{0}$, and let $\left(\tilde{\gamma}_{1}, \tilde{\gamma}_{2}, \alpha\right)$ be a geodesic hinge in $\tilde{S}$ with $\ell\left(\tilde{\gamma}_{i}\right)=\ell_{i}$. Then

$$
d_{S}\left(\gamma_{1}(0), \gamma_{2}\left(\ell_{2}\right)\right) \geq d_{\tilde{S}}\left(\tilde{\gamma}_{1}(0), \tilde{\gamma}_{2}\left(\ell_{2}\right)\right) .
$$

The essential idea here is that geodesics diverge more slowly in regions of higher curvature. Chavel gives hints for the solution of the exercise, but a complete demonstration of Lemma C. 3 may be found in [Kli95][Corollary 2.7.3]. As demonstrated by Klingenberg, Lemma C. 3 is a direct corollary of the Rauch comparison theorem.

A proper development of the Rauch theorem would require an exposition on Jacobi fields, which we will avoid here. A Jacobi field is a vector field defined on a geodesic. The field can be thought of as an indication of how a geodesic $\gamma$ would be displaced if its initial tangent direction were perturbed. Jacobi fields are defined as solutions to a differential equation which is reminiscent of the equation of a harmonic oscillator.

However, the essence of the Rauch theorem can be described without reference to Jacobi fields. In the words of Klingenberg [Kli95][p.215]:

We like to view Rauch's Theorem as a comparison between two infinitesimally slim triangles on manifolds $S, \tilde{S}$ with curvature satisfying $G \leq \tilde{G}$. The two triangles have two long sides and an infinitesimally small angle between them. There is a given correspondence between these two triangles such that the length of corresponding sides and the size of the included angle are the same. Then the length of the infinitesimal side opposite the infinitesimal angle in the triangle on $S$ is $\geq$ the infinitesimal side in the triangle on $\tilde{S}$. Corollary 2.7.3 [which implies Lemma C.3] is the integrated version of the Rauch theorem.

Thus the demonstration of Lemma C. 3 can be understood as applying the Rauch theorem to infinitesimal triangles $\left\{x, \gamma_{3}(t), \gamma_{3}(t+\delta)\right\}$, where $\gamma_{3}$ is a minimal geodesic between $\gamma_{2}\left(\ell_{2}\right)$
and $\gamma_{1}(0)$. Such a triangle is depicted in red in Figure C.2(b).
To finish the proof of Lemma C.1, Chavel refers to a result, [Cha06][Corollary IX.5.1], implying that Lemma C. 3 is equivalent to the assertion that the angles at $p$ and $q$ (angles $\alpha_{1}$ and $\alpha_{2}$ in Figure C.2(b)) are smaller than the corresponding angles at $\tilde{p}$ and $\tilde{q}$ in $\tilde{x} \tilde{p} \tilde{q}$. Let $\tilde{\gamma}_{3}$ be a minimal geodesic between $\tilde{p}$ and $\tilde{q}$ on $\tilde{S}$, and parameterize $\gamma_{3}$ and $\tilde{\gamma}_{3}$ from 0 at $p$ (respectively $\tilde{p}$ ) to 1 at $q$ (respectively $\tilde{q})$. Then a variational argument implies that for a very small $\delta, d_{S}\left(x, \gamma_{3}(\delta)\right)<d_{\tilde{S}}\left(\tilde{x}, \tilde{\gamma}_{3}(\delta)\right)$ and $d_{S}\left(x, \gamma_{3}(1-\delta)\right)<d_{\tilde{S}}\left(\tilde{x}, \tilde{\gamma}_{3}(1-\delta)\right)$. This result carries over to all parameter values by an infinitesimal argument, and the containment of $\gamma_{3}$ in $B_{S}(x ; \rho)$ follows because $\tilde{\gamma}_{3}$ is contained in $B_{\tilde{S}}(\tilde{x} ; \rho)$ on the sphere $\tilde{S}$.

## C. 3 Improving the bound in Theorem 5.19

The estimates on the injectivity radius and the scr in Section 5.3.1 can probably be tightened. In fact, our estimate on the injectivity radius, employed Lemma 5.16, which is a trivial consequence of Equation (5.7), the Fenchel-Reshetnyak Theorem applied to a loop. The Fenchel-Reshetnyak Theorem as given by Equation (5.7) states that the average curvature on $\gamma$ will be greater than $\frac{\pi}{\ell(\gamma)}$. It is clear that some points on $\gamma$ must have a curvature greater than this, otherwise $\gamma$ could not be a loop. In fact, we expect that just as for the case with smooth closed curves, the loop which minimizes the maximum curvature is in fact a circle. This means that Lemma 5.16 can most likely be strengthened to assert that:

$$
\begin{equation*}
\exists z \in \gamma \quad \text { such that } \quad \rho_{f}(z) \leq \frac{\ell(\gamma)}{2 \pi} . \tag{C.1}
\end{equation*}
$$

In this case, Theorem 5.17 could be improved to the more satisfactory

$$
\begin{equation*}
\rho_{i}(x) \geq\left(\frac{\pi}{1+\pi}\right) \rho_{f}(x) . \tag{C.2}
\end{equation*}
$$

Of course, the validity of Equation (C.1), would immediately imply an improvement on the estimate (5.9), which would improve the bound in Theorem 5.19 to Equation (5.8). However, we expect the bound can be improved even further.

We show that not only would Conjecture C. 4 below immediately imply Equation (C.2), but it would also be a more powerful tool for bounding the scr, because we can use it directly to tackle axioms 1 and 2 of strong convexity in the proof of Observation C. 6 below. In contrast, to tackle the first two axioms of strong convexity in the proof of Theorem 5.19, we employed a bound on the injectivity radius of each point. Thus with Conjecture C. 4
we can place the point witnessing the curvature bound directly on the violating geodesics, rather than simply somewhere within a disk containing the endpoints ( $p$ and $q$ ).

Conjecture C. 4 If distinct geodesics $\alpha$ and $\gamma$ connect $p$ and $q$ in $S$, then there is a point $z \in \alpha \cup \gamma$ with

$$
\rho_{f}(z) \leq \frac{\ell(\alpha)+\ell(\gamma)}{2 \pi} .
$$

The intuition behind this conjecture is that if two distinct geodesics connect $p$ and $q$, then there must be a witness to the necessary extrinsic curvature somewhere in the neighbourhood of the geodesics themselves. The claim immediately implies the better bound (C.2) on the injectivity radius.

Observation C. 5 Given the validity of Conjecture C.4, then for all $x \in S$,

$$
\rho_{i}(x) \geq\left(\frac{\pi}{1+\pi}\right) \rho_{f}(x) .
$$

Proof Suppose $r>\rho_{i}(x)$. Then there is a $p \in B_{S}(x ; r)$ such that two distinct geodesics, $\gamma$ and $\alpha$, connect $x$ with $p$ and are contained in $B_{S}(x ; r)$. By appealing to Klingenberg's Lemma 5.5, we may assume that $\ell(\gamma)+\ell(\alpha) \leq 2 r$. Then Conjecture C. 4 ensures that there is a $z \in \gamma \cup \alpha$ such that

$$
\rho_{f}(z) \leq \frac{\ell(\gamma)+\ell(\alpha)}{2 \pi} \leq \frac{r}{\pi},
$$

and by the Lipschitz continuity of the lfs,

$$
\rho_{f}(x) \leq \rho_{f}(z)+d_{S}(x, z) \leq\left(\frac{1+\pi}{\pi}\right) r .
$$

As well, we get a much more satisfactory bound on the scr.
Observation C. 6 Given the validity of Conjecture C.4, then for all $x \in S$,

$$
\rho_{\mathrm{sc}}(x) \geq \frac{1}{2}\left(\frac{\pi}{1+\pi}\right) \rho_{f}(x) .
$$

The demonstration of Observation C. 6 can be done without the aid of Lemmas 5.15 or 5.16 , but we do require an additional elementary observation on curves and surfaces. Also, we still require Chavel's Lemma 5.18, which employs more sophisticated Riemannian geometry, as described in Section C.2.

Lemma C. 7 ([Top06][p.55]) Suppose $\alpha$ is a convex planar curve connecting $p_{\alpha}$ and $q_{\alpha}$, and suppose $\gamma$ is an arbitrary space curve with endpoints $p$ and $q$ and $\ell=\ell(\alpha)=\ell(\gamma)$. Let both curves be parameterized by the arclength parameter $t$ and denote the curvatures of $\alpha(t)$ and $\gamma(t)$ by $k_{\alpha}(t)$ and $k_{\gamma}(t)$ respectively.

If $k_{\alpha}(t) \geq k_{\gamma}(t)$ for all $t \in[0, \ell]$ then

$$
d_{\mathbb{R}^{3}}\left(p_{\alpha}, q_{\alpha}\right) \leq d_{\mathbb{R}^{3}}(p, q),
$$

and equality holds if and only if $\alpha$ and $\gamma$ are identified by a rigid motion of space.
Toponogov's observation, Lemma C.7, yields a corresponding result for geodesics:
Lemma C. 8 If $\gamma$ is a geodesic contained in a Euclidean ball of radius $\rho$, then either $\ell(\gamma)<$ $\pi \rho$, or there exists a point $z \in \gamma$ where $\rho_{\kappa}(z)<\rho$.

Proof Assume $\gamma$ connects $p$ and $q$. If $\ell(\gamma)>\pi \rho$, consider the subsegment $\hat{\gamma}$ connecting $p$ with $\hat{q}$, such that $\ell(\hat{\gamma})=\pi \rho$. Then if $\rho_{\kappa} \geq \rho$ on $\hat{\gamma}$, we compare $\hat{\gamma}$ with a semicircular arc of radius $\rho$. Since $k_{\gamma}(t) \leq \kappa(\gamma(t)) \leq \frac{1}{\rho}$, Toponogov's Lemma C. 7 implies that $d_{\mathbb{R}^{3}}(p, \hat{q}) \geq 2 \rho$, the distance between the endpoints of the semicircular arc. But this means that $p$ and $\hat{q}$ cannot both be contained in an open Euclidean ball of radius $\rho$, a contradiction.

The proof of Observation C. 6 is a modification of part of the proof of Theorem 5.19. For convenience we give the complete proof including those parts that are unchanged.

Proof of Observation C. 6 Consider the geodesic disk of radius $r$ centred at $x$. There are three ways in which $B_{S}(x ; r)$ can fail to be strongly convex: There exist $p, q \in B_{S}(x ; r)$ such that either
(i) the minimizer $\gamma$ connecting $p$ and $q$ is not unique, or
(ii) in addition to $\gamma$, there is another geodesic $\alpha$ connecting $p$ and $q$ and contained in

$$
B_{S}(x ; r) \text {, or }
$$

(iii) $\gamma$ is not contained in $B_{S}(x ; r)$.

For case (i), suppose $\alpha$ is another minimizer connecting $p$ and $q$. Since $d_{S}(p, q)<2 r$, by Conjecture C. 4 there is a $z \in \alpha \cup \gamma$ with $\rho_{f}(z) \leq \frac{2 r}{\pi}$. Without loss of generality, assume
$d_{S}(p, z) \leq d_{S}(q, z)$. Then by the triangle inequality,

$$
\begin{aligned}
d_{S}(x, z) & \leq d_{S}(x, p)+d_{S}(p, z) \\
& \leq 2 r .
\end{aligned}
$$

Therefore, by the Lipschitz continuity of the lfs,

$$
\rho_{f}(x) \leq \rho_{f}(z)+2 r \leq\left(\frac{2+2 \pi}{\pi}\right) r,
$$

and so

$$
\begin{equation*}
r \geq \frac{1}{2}\left(\frac{\pi}{1+\pi}\right) \rho_{f}(x) . \tag{C.3}
\end{equation*}
$$

For case (ii), $\alpha$ is not necessarily minimal, but it is contained in $B_{S}(x ; r)$. By Lemma C.8, either
(a) $\ell(\alpha) \leq 2 \pi r$, or
(b) $\exists z \in \alpha$ with $\rho_{\kappa}(z)<r$.

In case (a), $\ell(\alpha)+\ell(\beta) \leq 2 \pi r+2 r$, and Conjecture C. 4 implies that there is a $z \in \alpha \cup \gamma \subset$ $B_{S}(x ; r)$ with

$$
\rho_{f}(z) \leq \frac{(2(\pi+1) r}{2 \pi}=\left(\frac{\pi+1}{\pi}\right) r .
$$

By the Lipschitz continuity of the lfs, $\rho_{f}(x) \leq \rho_{f}(z)+r=\left(\frac{2 \pi+1}{\pi}\right) r$, and so

$$
\begin{equation*}
r \geq\left(\frac{\pi}{2 \pi+1}\right) \rho_{f}(x) \tag{C.4}
\end{equation*}
$$

In case (b), Lipschitz continuity of the lfs gives $\rho_{f}(x) \leq \rho_{f}(z)+d_{S}(x, z) \leq 2 r$, and so

$$
\begin{equation*}
r \geq \frac{1}{2} \rho_{f}(x) . \tag{C.5}
\end{equation*}
$$

Equation (C.3) provides the smallest of the three bounds (C.3), (C.4), and (C.5), so provided

$$
\begin{equation*}
r<\frac{1}{2}\left(\frac{\pi}{\pi+1}\right) \rho_{f}(x), \tag{C.6}
\end{equation*}
$$

cases (i) and (ii) cannot happen.
It remains to consider case (iii). For this we assume that Equation (C.6) is satisfied and we turn to Chavel's Lemma 5.18. Since Equation (C.6) is satisfied we have $\rho_{i}(x)>2 r$, by Observation C.5.

Consider the geodesic triangle $x p q$ consisting of minimal geodesics, with notation as in Lemma 5.18. By hypothesis we now have $\rho_{i}(x) \geq 2 r$, and also $\sum_{i=1}^{3} \ell\left(\gamma_{i}\right)<4 r$. Thus the conditions of Lemma 5.18 are satisfied with $\rho_{0}=2 r$, provided there is no $z \in B_{S}(x ; 2 r)$ with $2 \pi \rho_{G}(z)<4 r$. If there were such a $z$, it would imply $\rho_{f}(z)<\frac{2 r}{\pi}$, and the Lipschitz continuity of the local feature size yields

$$
\begin{equation*}
r>\frac{1}{2}\left(\frac{\pi}{1+\pi}\right) \rho_{f}(x) \tag{C.7}
\end{equation*}
$$

contradicting the hypothesis that Equation (C.6) is satisfied, and thus proving the claim.

The improved bound in Observation C. 6 gives a relaxation in the lfs-based extrinsic sampling radius needed to satisfy the intrinsic sampling radius. Using $C=\frac{1}{2}\left(\frac{\pi}{\pi+1}\right)$ in Equation (5.13) reveals that $\epsilon_{f} \leq 0.298$ would be sufficient to satisfy the criteria of Theorem 5.14.

## Appendix D

## Allowable region calculations

In this appendix we consider the problem where we are given three vertices $u$, $v$, and $p$, fixed in space, and we wish to know where in space we can place a fourth vertex $w$ such that edge $[p, w]$ is locally Delaunay in hinge $(u,[p, w], v)$. We refer to the set of locations where $[p, w]$ is locally Delaunay as the feasible region for $w$.

In Section D. 1 we make some general observations about the feasible region and its boundary, and establish some of the tools that will facilitate the calculations of Section D.2, where we set out to define a set of half spaces whose intersection is completely contained within the feasible region of $w$. We refer to the region thus delimited as the allowable region for $w$. The primary motivation for the description of the allowable region is the decimation algorithm of Section 6.3.

## D. 1 The feasible region

We are interested in the feasible region for $w$ in hinge $(u,[p, w], v)$. Let $\theta_{1}(w)=\angle w u p$, $\theta_{2}(w)=\angle w v p$, and define

$$
\begin{equation*}
f(w)=\theta_{1}+\theta_{2} . \tag{D.1}
\end{equation*}
$$

Then $w$ is a feasible point if $f(w) \leq \pi$. An equivalent way of characterizing the feasible region is by means of the cotan formula. Define

$$
\begin{equation*}
g(w)=\cot \theta_{1}+\cot \theta_{2} . \tag{D.2}
\end{equation*}
$$

Then $w$ is a feasible point if $g(w) \geq 0$. The equivalence is easily checked by expanding the cotangents in terms of cosine and sine, and using the angle sum formulas. The feasible region


Figure D.1: The feasible region in the plane. Elementary considerations reveal that the region shaded pink lies within the feasible region, and the region shaded cyan lies outside of the feasible region. The boundary of the feasible region lies somewhere in the region shaded yellow.
includes its boundary, which we denote by $F=f^{-1}(\pi)=g^{-1}(0)$, but when $w$ lies on $F$, $(u,[p, w], v)$ is not in general position. The functions $f$ and $g$ are not well defined at $u$ and $v$, but it is natural to consider these points to lie on $F$. Therefore, we define $f(u)=f(v)=\pi$, and therefore $g(u)=g(v)=0$. The boundary surface $F$ is smooth everywhere except at $u$ and $v$, where it is not even continuous. Indeed $f$ itself is smooth over all space except at these two points.

For the purposes of exposition we refer to $\operatorname{aff}([u, p, v])$ as the horizontal plane, and lines and planes which are perpendicular to it are vertical. The circumcircle of $[u, p, v]$ is denoted $C$.

## D.1.1 Feasible points on the horizontal plane

As a first step towards describing $F$, we consider $w$ to be restricted to the horizontal plane. We know from the theory of planar Delaunay triangulations that if $w$ lies in the wedge defined by $\angle u p v$, then $w$ is a feasible point if and only if it lies on or within $C$. However,
we do not want to restrict our attention to locations for which $[u, p, w]$ and $[w, p, v]$ form a valid planar triangulation; $f(w)$ is defined over the entire plane.

Referring to Figure D.1, we identify $p^{*}$, the point antipodal to $p$ in $C$, as a significant point for our considerations. Consider the dashed line through $v$ and $p^{*}$. This line is perpendicular to $[p, v]$, and hence if $w$ lies on the same side of this line as $p$, then $\theta_{2}$ is necessarily acute, otherwise it is obtuse. Likewise, the dashed line through $u$ and $p^{*}$ defines the boundary between the region where $\theta_{1}(w)$ is acute, and the region where it is obtuse.

Thus the intersection of the two half-spaces where $\theta_{1}$ and $\theta_{2}$ are both acute, must necessarily lie within the feasible region. Likewise, the region where both $\theta_{1}$ and $\theta_{2}$ are obtuse, must lie outside of the feasible region.

These observations are illustrated in Figure D.1.

## D.1.2 The Lifting Lemma

We now consider points that are vertically displaced from the horizontal plane. Since our goal in Section D. 2 will be to establish vertical planes which bound an allowable region within the feasible region, we are especially interested in deciding when a point $q$ that lies in the horizontal plane is such that the entire vertical line through $q$ lies within the feasible region.

For $i=1,2$, if $\theta_{i}(q)$ is acute (obtuse), then $\theta_{i}(w)$ will also be acute (obtuse) for any $w$ on the vertical line through $q$. Indeed, the plane perpendicular to $[u, p]$ separates the region where $\theta_{1}$ is acute from the region where it obtuse, just as the corresponding line did in the horizontal plane. Therefore, if $\theta_{1}(q)$ and $\theta_{2}(q)$ are both acute, then $f(w)<\pi$ for any $w$ on the vertical line through $q$. Likewise, if both angles are obtuse, then the vertical line through $q$ will not intersect the feasible region.

Thus we are interested in the case where one of the angles is obtuse, and the other is acute. Let $\theta_{a}$ be the acute angle and $\theta_{o}$ the obtuse one. In this context a convenient alternate expression for the the cotan formula is

$$
\begin{equation*}
g(w)=\tan \left(\pi / 2-\theta_{a}(w)\right)-\tan \left(\theta_{o}(w)-\pi / 2\right) . \tag{D.3}
\end{equation*}
$$

Consider a point $q$ in the horizontal plane, and for the purposes of this exposition, assume that $\theta_{1}(q)$ is acute and $\theta_{2}(q)$ is obtuse. Let $P_{a}$ be the plane perpendicular to $[p, u]$, and $P_{o}$ the one perpendicular to $[p, v]$. As shown in Figure D.2(a), let $x_{a}$ and $x_{o}$ be the orthogonal projection of $q$ onto $P_{a}$ and $P_{o}$ respectively, and let $a_{a}=\left|\left[q, x_{a}\right]\right|, a_{o}=\left|\left[q, x_{o}\right]\right|$,


Figure D.2: Diagram for the Lifting Lemma calculations. (a) shows the planar quad with vertex $q$, (b) shows the hinge ( $u,[p, w], v$ ), with $w$ vertically above $q$. Elements that lie in $\operatorname{aff}([w, p, v])$ are drawn in red.
$b_{a}=\left|\left[p, x_{a}\right]\right|$, and $b_{o}=\left|\left[p, x_{o}\right]\right|$. Finally, let $\alpha_{a}=\pi / 2-\theta_{a}(q)$, and $\alpha_{o}=\theta_{o}(q)-\pi / 2$. We can now express Equation (D.3) for point $q$ as

$$
\begin{equation*}
g(q)=\frac{a_{a}}{b_{a}}-\frac{a_{o}}{b_{o}} . \tag{D.4}
\end{equation*}
$$

We now consider point $w$ on a vertical line through $q$. We construct the analogous elements as above. Thus let $\tilde{x}_{a}$ and $\tilde{x}_{o}$ be the orthogonal projection of $w$ onto $P_{a}$ and $P_{o}$ respectively, and let $\tilde{a}_{a}=\left|\left[w, \tilde{x}_{a}\right]\right|, \tilde{a}_{o}=\left|\left[w, \tilde{x}_{o}\right]\right|, \tilde{b}_{a}=\left|\left[p, \tilde{x}_{a}\right]\right|, \tilde{b}_{o}=\left|\left[p, \tilde{x}_{o}\right]\right|, \tilde{\alpha}_{a}=$ $\pi / 2-\theta_{a}(w)$, and $\tilde{\alpha}_{o}=\theta_{o}(w)-\pi / 2$. The quantities associated with the obtuse angle are shown in Figure D.2(b) for example.

Since $\left[x_{o}, \tilde{x}_{o}\right]$ is parallel to $[q, w]$, we have $\tilde{a}_{o}=a_{o}$. Likewise, $\tilde{a}_{a}=a_{a}$. Let $h=|[q, w]|$. Then $\tilde{b}_{o}=\left(b_{o}^{2}+h^{2}\right)^{1 / 2}$ and $\tilde{b}_{a}=\left(b_{a}^{2}+h^{2}\right)^{1 / 2}$, and we can express Equation (D.3) for $w$ as

$$
\begin{equation*}
g(w)=\frac{a_{a}}{\left(b_{a}^{2}+h^{2}\right)^{1 / 2}}-\frac{a_{o}}{\left(b_{o}^{2}+h^{2}\right)^{1 / 2}} . \tag{D.5}
\end{equation*}
$$

Point $w$ lies in the feasible region if and only if $g(w) \geq 0$. Solving this equation for $h$ yields

$$
\left(a_{o}^{2}-a_{a}^{2}\right) h^{2} \leq a_{a}^{2} b_{o}^{2}-a_{o}^{2} b_{a}^{2},
$$

and if $a_{a}<a_{o}$, we require

$$
\begin{equation*}
h^{2} \leq \frac{a_{a}^{2} b_{o}^{2}-a_{o}^{2} b_{a}^{2}}{a_{o}^{2}-a_{a}^{2}} . \tag{D.6}
\end{equation*}
$$

If $q$ lies in the feasible region, then Equation (D.4) reveals that the right-hand side of Equation (D.6) is positive, and thus there is a maximum value for $h$, beyond which $w$ will lie outside of the feasible region.

On the other hand, if $a_{a}>a_{o}$, then we obtain

$$
\begin{equation*}
h^{2} \geq \frac{a_{a}^{2} b_{o}^{2}-a_{o}^{2} b_{a}^{2}}{a_{o}^{2}-a_{a}^{2}}, \tag{D.7}
\end{equation*}
$$

and since the right-hand side is now non-positive if $q$ lies in the feasible region, then $w$ will also be a feasible point, regardless of the value of $h$. Observe also that if $a_{a}=a_{o}$, then from Equation (D.5), $w$ will be a feasible point if and only if $b_{a} \leq b_{o}$, which will be the case if and only if $q$ is a feasible point.

The primary observation of interest for our construction of the allowable region in Section D. 2 is summarized by:

Lemma D. 1 (Lifting) Let $q$ be a feasible point in the horizontal plane. If one of $\theta_{1}(q)$, $\theta_{2}(q)$ is obtuse, then the vertical line through $q$ lies within the feasible region if and only if $a_{a} \geq a_{o}$.

## D.1.3 Feasible points on the circumcylinder

It is useful to consider the case where $w$ is restricted to the circumcylinder of $[u, p, v]$ : the right circular cylinder defined by $C$. If $q$ lies on the arc of $C$ which subtends $\angle u p v$, then the numerator in Equations (D.6) and (D.7) vanishes. Thus in this case, if $a_{a}<a_{o}, q$ is the only point on the vertical line through $q$ which lies in the feasible region.

For all points on the arc subtending $\angle u p v$, we have $\alpha_{a}=\alpha_{o}$, since $f(q)=\pi$ for such points. Suppose for the moment that $\theta_{a}=\theta_{1}$. The inequality $a_{a} \geq a_{o}$, can be related to an inequality in the lengths of the edges $[u, q]$ and $[v, q]$ by the relations $a_{a}=|[u, q]| \sin \alpha$, and $a_{o}=|[v, q]| \sin \alpha$, where $\alpha=\alpha_{a}$. Since $\sin \alpha>0, a_{a} \geq a_{o}$ is equivalent to $|[u, q]| \geq|[v, q]|$.

We obtain
Lemma D. 2 (Circumcylinder) Let $q$ be a point on the arc of $C$ that subtends $\angle u p v$, and let $w$ lie on a vertical line through $q$. Then $w$ is a feasible point if and only if $\theta_{1}$ is obtuse and $|[u, q]| \leq|[v, q]|$ or $\theta_{1}$ is acute and $|[u, q]| \geq|[v, q]|$.


Figure D.3: Feasible regions for points on the circumcylinder of $[u, p, v]$. Point $p^{*}$ is antipodal to $p$ in the circumcircle, and $m$ is the midpoint of the arc subtending $\angle u p v$. Lemma D. 2 asserts that if $w$ lies vertically above or below one of the red arcs, then $[p, w]$ will be nlD in hinge $(u,[p, w], v)$. The green line delineates the points where $a_{a}=a_{o}$, as described in the text.

We are naturally lead to consider Lemma D. 2 in two cases, according to whether $u$ and $v$ lie on the same side of $\left[p, p^{*}\right]$ or not, where $p^{*}$ is the point antipodal to $p$, as depicted in Figure D.3. Let $m$ be the midpoint of the arc subtending $\angle u p v$. Then since $|[u, m]|=|[v, m]|$, it follows that $a_{a}=a_{o}$ and the vertical line through $m$ always lies in the feasible region.

In the opposite side case, Figure D.3(a), if a point $w$ is located vertically above or below the arc $p^{*} m(\mathrm{red})$, then $[p, w]$ is nlD in hinge $(u,[p, w], v)$. However if $w$ is located anywhere vertically above or below arcs $\overparen{m u}$ or $\widehat{v p^{*}}$, then $[p, w]$ will be locally Delaunay. In the sameside case, Figure D.3(b), $[p, w]$ is nlD for $w$ anywhere above or below the arc $\overparen{m u}$, where $u$ is the vertex closest to $p^{*}$.

The dashed lines in Figure D. 3 are perpendicular to $[p, u]$ and $[p, v]$. As discussed in Section D.1.1, the perpendicular through $u$, divides the plane according to a half-space
where $\theta_{1}$ is acute and one where $\theta_{1}$ is obtuse, and the perpendicular through $v$ does likewise for $\theta_{2}$. From consideration of these spaces we see that in the opposite-side case, all points on the circumcylinder that are not on vertical lines through the (red) arc $p^{*} m$, lie in the feasible region. In the same-side case, points that lie on vertical lines through the (blue) arc $\widehat{u p}^{*}$ are outside of the feasible region, but all other points that are not specifically addressed by Lemma D. 2 are feasible.

The set of points in the horizontal plane for which $a_{a}=a_{o}$ is depicted with the green line in Figure D.3. This line contains $m$ and $p^{*}$, and it bisects an angle $\gamma$ formed by the dashed perpendicular lines. This angle is $\gamma$ and it is equal to $\angle u p v$. This (green) line divides the plane into two half-spaces. If $q$ lies on the side which contains $p$, then either $\theta_{1}$ and $\theta_{2}$ are both acute, or $a_{a}>a_{o}$. Thus by the lifting lemma, a vertical line through a feasible point in this region will lie entirely within the feasible region.

## D. 2 Linear allowable region boundaries

Given hinge $(u,[p, w], v)$, we again consider $u, p$, and $v$ to be fixed, but now consider the feasible region of $w$ over all space. In this section our goal is to identify a set of half-spaces whose intersection defines an allowable region for $w$. The strategy will be to establish the appropriate lines in the horizontal plane, and use the Lifting Lemma D. 1 to verify the correctness of the vertical constraint planes through these lines.

A constraint plane is correct if the part of it that forms the boundary of the allowable region is completely contained within the feasible region. Indeed, on any ray emanating from $p, f(w)$ is monotonically increasing. Therefore, an allowable region which contains $p$ and whose boundary does not intersect $F$, must necessarily be completely contained in the feasible region.

As indicated in Figure D.3, a natural choice for correct constraint planes would be the planes perpendicular to edges $[p, u]$ and $[p, v]$. These planes delimit an allowable region where both $\theta_{1}$ and $\theta_{2}$ are constrained to be acute. However, for the purposes of the decimation algorithm of Section 6.3, we found these constraint planes to be too restrictive. Experiments revealed cases where clusters of vertices around sharp ridges would not be decimated.

Instead, we choose two constraint planes which maximize the volume of the allowable region within the circumcylinder of $[u, p, v]$. In some cases we will require a third plane to further contain the allowable region outside of the circumcylinder. As depicted in Figure D.3,


Figure D.4: Opposite-side case
we consider two distinct cases, according to whether or not $u$ and $v$ lie on the same-side of the diameter through $p$.

## D.2.1 Opposite-side case

The point $m$, which is the midpoint of the arc of $C$ subtending $\angle u p v$, maximizes the area of the triangle $[u, v, m]$. We construct two vertical constraint planes which contain the segments $[m, u]$ and $[m, v]$ respectively. We refer to Figure D.4, where these segments are contained in the red lines.

Within the circumcylinder, these planes are contained in the feasible region. Indeed, the standard arguments for planar Delaunay triangulations ensures that $f(q) \leq \pi$ for any $q$ on $[m, u]$ or $[m, v]$, and the Lifting Lemma D. 1 applies for all these points.

Now consider the portions of these constraint planes which lie outside of the circumcylinder and potentially form the boundary of the allowable region. In the plane of Figure D.4, we are considering the red lines that extend through $u$ and $v$.

Since $v$ is closer to $p^{*}$ than $u$ in Figure D.4, we refer to it as the close vertex, the
other vertex $(u)$ is the far vertex. Since $m$ is closer to the far vertex than is $p^{*}$, the line containing $m$ and $u$ makes a smaller angle with the tangent to the circle at $u$ than does the line containing $p^{*}$ and $u$. Therefore, outside of the circumcylinder, the constraint plane through the far vertex is contained in the feasible region because both $\theta_{1}$ and $\theta_{2}$ must be acute there.

On the other hand, the constraint plane containing the close vertex may intersect $F$. Without loss of generality, we will assume that $v$ is the close vertex, as depicted in Figure D.4. Although $f$ is not smooth at $v, F$ constrained to aff $([u, p, v])$ has an "external" tangent vector at $v$. Let $\varphi=\angle v u p$. Then a line through $v$ and containing a point $q$ outside the circle will have points close to $v$ which are outside the feasible region if $\angle q v p>\pi-\varphi$. We need to ensure that our constraint plane lies inside this angle.

Let $\ell$ be the half-line terminating at $m$ and extending through $v$, and let $\alpha=\angle m v p^{*}$, and $\gamma=\angle u p v$. For a point $q$ on $\ell$, outside of $C$, we have $\angle q v p=\pi / 2+\alpha$. We need to ensure that

$$
\begin{equation*}
\pi / 2+\alpha+\varphi \leq \pi \tag{D.8}
\end{equation*}
$$

Since $\widehat{v m}$ is half the length of $\widehat{v u}, \angle m u v=\gamma / 2$. Thus, referring to $\angle m u p$ in Figure D.4, we have

$$
\pi / 2+\alpha=\varphi+\gamma / 2
$$

Plugging into Equation (D.8), we get the requirement that

$$
\begin{equation*}
\alpha \leq \gamma / 4 \tag{D.9}
\end{equation*}
$$

Since $\alpha$ is subtended by $p^{*} m$, and $p^{*}$ lies in $\overparen{v m}$ which subtends $\gamma / 2$, it follows that Equation (D.9) will be satisfied if and only if

$$
\left|\left[p^{*}, m\right]\right| \leq\left|\left[v, p^{*}\right]\right| .
$$

If this is not the case, we add another vertical constraint plane which goes through $v$ and the midpoint of $\overparen{v m}$ so that equality is attained in the constraint (D.9).

We now verify that $\ell$ lies entirely within the feasible region when Equation (D.9) is satisfied. Suppose that aff $([u, p])$ does not intersect $\ell$. For $q$ on $\ell$, external to the circle, $\angle q v p=\pi / 2+\alpha$ and $\angle q u p \leq \varphi$. It follows then from Equation (D.8) that $q$ is in the feasible region.


Figure D.5: Comparing tangents

Now suppose that aff $([u, p])$ intersects $\ell$ at $z$. If $q$ lies between $v$ and $z$, then the above argument still applies. Suppose then that $z$ lies between $q$ and $v$. In this case the angle $\angle q u p$ increases as $|[v, q]|$ increases. However, by considering the triangle $[v, u, q]$, we see that

$$
\angle v u q+\angle u v q<\pi
$$

for any $q \in \ell$. Since $\angle v u q=\theta_{1}+\varphi$, and $\angle u v q=\theta_{2}+\angle p v u$, it follows that $f(q)<\pi$, so $q$ is in the feasible region.

It remains to verify that the Lifting Lemma D. 1 applies on $\ell$, ensuring that the vertical plane through $\ell$ is entirely contained in the feasible region. Referring to Figure D.5, we have $a_{a}=|[u, q]| \sin \alpha_{a}, a_{o}=|[v, q]| \sin \alpha_{o}$, and $|[u, q]|>|[v, q]|$, because $[u, q]$ is opposite the obtuse angle $\angle u v q$ in triangle $[u, q, v]$. Therefore, since $\pi / 2 \geq \alpha_{a}>\alpha=\alpha_{o}$, by our construction of $\ell$, we have $a_{a}>a_{o}$ and the Lifting Lemma D. 1 applies.


Figure D.6: Same-side case

## D.2.2 Same-side case

It is evident from Figure D.3(b), that the constraint planes constructed for the opposite-side case will not suffice for the same-side case. We will assume that $u$ is the close vertex (closer to $\left.p^{*}\right)$ and $v$ is the far vertex, as depicted in Figure D.3(b). Since points on the circumcylinder above or below arcs $\overparen{m u}$ (red) and $\widehat{u p^{*}}$ (blue) are outside of the feasible region, the entire arc $\overparen{m p^{*}}$ must be excluded from the allowable region. We therefore construct a vertical constraint plane that contains $\left[m, p^{*}\right]$. We also construct a vertical constraint plane that includes $m$ and the far vertex. The correctness of this latter plane is verified as in the opposite-side case.

In order to evaluate the constraint plane through $p^{*}$, we will again invoke the formulation (D.3). Let $\ell$ be the half-line starting at $m$ and containing $p^{*}$, and let $q$ be a point on $\ell$ that lies within the circumcircle. Again let $\gamma$ be the angle subtended by arc $\widehat{v u}$, and notice that $\ell$ bisects $\angle u p^{*} v=\gamma$, by construction of $m$.

Let $y$ be the point of intersection between $[p, u]$ and $\ell$. If $q$ lies between $m$ and $y$, then $q$ is
in the feasible region by standard arguments for planar Delaunay triangulations. Therefore, assume $q$ lies between $y$ and $p^{*}$, as shown in Figure D.6.

Let $a=a_{a}=a_{o}$. Then

$$
g(q)=\frac{a}{b_{a}}-\frac{a}{b_{o}} .
$$

But since $b_{a}=\left|\left[p^{*}, u\right]\right|-\left|\left[p^{*}, x_{a}\right]\right|$ and $b_{o}=\left|\left[p^{*}, v\right]\right|-\left|\left[p^{*}, x_{o}\right]\right|$, and $\left|\left[p^{*}, x_{a}\right]\right|=\left|\left[p^{*}, x_{o}\right]\right|$, we have $b_{a}<b_{o}$, because $u$ is the close point. It follows that $g(q)>0$, and therefore $q$ lies in the feasible region. It follows directly from the Lifting Lemma D. 1 that any point $w$ vertically displaced from $q$ is also in the feasible region, and thus the restriction of vertical constraint plane through $\ell$ to the circumcylinder lies entirely within the feasible region.

Similar arguments reveal that $\ell$ leaves the feasible region when it exits the circumcylinder. The restriction of $F$ to $\operatorname{aff}([u, p, v])$ crosses $\ell$ at $p^{*}$ and does not intersect it at any other point, but $\ell$ does serve as the asymptote for $F$. Indeed, for $q \in \ell$, we have $\theta_{1} \rightarrow \pi / 2+\gamma / 2$ and $\theta_{2} \rightarrow \pi / 2-\gamma / 2$, as $|[q, p]| \rightarrow \infty$.

Anyway, we need a new vertical constraint plane in order to create a valid boundary for the allowable region exterior to the circumcylinder and on the other side of $\operatorname{aff}\left(\left[p, p^{*}\right]\right)$ than $u$ and $v$. The best plane we can choose, in terms of maximizing the area of the allowable region, is the one containing the tangent to $F \cap \operatorname{aff}([u, p, v])$ at $p^{*}$.

Note that whereas the obtuse angle was at the far vertex inside the circumcylinder, in the region containing $F$ outside of the circumcylinder, the obtuse angle is at the close vertex. Set
$k=\frac{\left.\left\|v, p^{*}\right\|\right]}{\left[u, p^{*}\right]}$, and let $\ell_{t}$ be the half-line starting at $p^{*}$ and extending away from the circle such that for any $q \in \ell_{t}$, we have $\frac{a_{a}}{a_{o}}=k$, where $a_{o}$, and $a_{a}$ are the distances of $q$ from aff $\left(\left[u, p^{*}\right]\right)$ and $\operatorname{aff}\left(\left[v, p^{*}\right]\right)$ respectively, as shown in Figure D.7. Let $\tau$ be the angle $\ell_{t}$ makes with $\operatorname{aff}\left(\left[v, p^{*}\right]\right)$. Putting $s=\left|\left[p^{*}, q\right]\right|$, we get

$$
s=\frac{a_{2}}{\sin \tau}=\frac{a_{1}}{\sin (\gamma-\tau)},
$$

or

$$
\frac{\sin \tau}{\sin (\gamma-\tau)}=k
$$



Figure D.7: Exit tangent

Expanding the denominator we get

$$
\begin{equation*}
\tan \tau=\frac{k \sin \gamma}{(1+k \cos \gamma)} \tag{D.10}
\end{equation*}
$$

We have $b_{o}=\left|\left[u, x_{o}\right]\right|$, and $b_{a}=\left|\left[v, x_{a}\right]\right|$. Set $b=\left|\left[u, p^{*}\right]\right|$, so $b_{o}=b+s \cos (\gamma-\tau)$ and $b_{a}=b k+s \cos \tau$. Using $a_{a}=s \sin \tau$ and $a_{o}=\left(\frac{\sin \tau}{k}\right) s$, we get

$$
\begin{equation*}
g(q)=\frac{a_{a}}{b_{a}}-\frac{a_{o}}{b_{o}}=\frac{s \sin \tau}{b k+s \cos \tau}-\frac{s \sin \tau}{b k+s k \cos (\gamma-\tau)} . \tag{D.11}
\end{equation*}
$$

By our construction of $\ell_{t}, \pi / 2>\tau>\gamma-\tau$. Thus $\cos (\gamma-\tau)>\cos \tau$. Also $k>1$, because $v$ is the far vertex. It follows then that $g(q)>0$ for $s>0$, and thus $q$ lies in the feasible region. Also, we have $a_{a}>a_{o}$ by construction, so the Lifting Lemma D. 1 ensures the validity of the vertical constraint plane containing $\ell_{t}$.

Finally, differentiating (D.11) with respect to $s$, the variable that parameterizes $\ell_{t}$, we get

$$
\left.\frac{d g}{d s}\right|_{s=0}=0
$$

confirming that $\ell_{t}$ is indeed tangent to $F$ at $p^{*}$, as desired.

## Bibliography

[AAH02] O. Aichholzer, L.S. Alboul, and F. Hurtado, On flips in polyhedral surfaces, International Journal of Foundations of Computer Science 13 (2002), no. 2, 303-311.
[AB98] Nina Amenta and Marshall W. Bern, Surface reconstruction by Voronoi filtering, Symp. Comp. Geom., 1998, pp. 39-48.
[ABE98] Nina Amenta, Marshall Bern, and David Eppstein, The crust and the betaskeleton: combinatorial curve reconstruction, Graphical Models and Image Processing 60 (1998), no. 2, 125-135.
[ABK98] Nina Amenta, Marshall Bern, and Manolis Kamvysselis, A new voronoi-based surface reconstruction algorithm, ACM SIGGRAPH, 1998.
[ABR06] Lyuba Alboul, Willie Brink, and Marcos Rodrigues, Mesh optimisation based on Willmore energy, Euro. Workshop on Comp. Geom., 2006, pp. 133-136.
[ACDL00] Nina Amenta, Sunghee Choi, Tamal K. Dey, and Naveen Leekha, A Simple Algorithm for Homeomorphic Surface Reconstruction, Symp. Comp. Geom., 2000, pp. 213-222.
[AD07] N. Amenta and T. Dey, Normal variation for adaptive feature size, Erratum, 2007.
[AF02] Ernst Althaus and Christian Fink, A polyhedral approach to surface reconstruction from planar contours, Integer Programming and Combinatorial Optimization, 2002, pp. 258-272.
[AGJ00] Udo Adamy, Joachim Giesen, and Matthias John, New techniques for topologically correct surface reconstruction, IEEE Visualization, 2000, pp. 373-380.
[AGMR98] G. Albers, L.J. Guibas, J.S.B. Mitchell, and T. Roos, Voronoi diagrams of moving points, International Journal of Computational Geometry and Applications 8 (1998), no. 3, 365-380.
[AK00] Franz Aurenhammer and Rolf Klein, Voronoi diagrams, Handbook of Computational Geometry (J. Sack and G. Urrutia, eds.), Elsevier Science Publishing, 2000, pp. 201-290.
[AKTvD00] Lyuba Alboul, Gertjan Kloosterman, Cornelis Traas, and Ruud van Damme, Best data-dependent triangulations, J. Comput. Appl. Math. 119 (2000), no. 12, 1-12.
[Alb03] L. Alboul, Optimising triangulated polyhedral surfaces with self-intersections, 10th IMA international conference on the Mathematics of Surfaces, SpringerVerlag, 2003, Lecture notes in computer science, pp. 48-72.
[AM01] Ernst Althaus and Kurt Mehlhorn, Traveling salesman-based curve reconstruction in polynomial time, SIAM J. Comput. 31 (2001), no. 1, 27-66.
[APR03] Nina Amenta, Thomas J. Peters, and Er Russell, Computational topology: ambient isotopic approximation of 2-manifolds, Theoretical Computer Science 305 (2003), 3-15.
[AS00] Marco Attene and Michela Spagnuolo, Automatic surface reconstruction from point sets in space, Computer Graphics Forum 19 (2000), 457-465.
[ASY08] Pankaj K. Agarwal, Bardia Sadri, and Hai Yu, Untangling triangulations through local explorations, Symp. Comp. Geom., 2008, pp. 288-297.
[AT72] Louis I. Alpert and L.V. Toralballa, An elementary definition of surface area in $\mathbf{E}^{n+1}$ for smooth surfaces, Pacific Journal of Mathematics 40 (1972), no. 2, 261-268.
[AZ67] A. D. Aleksandrov and V. A. Zalgaller, Intrinsic geometry of surfaces, Transactions of Mathematical Monographs, vol. 15, AMS, 1967.
[BA05] J. Andreas Baerentzen and Henrik Aanaes, Signed distance computation using the angle weighted pseudonormal, IEEE Transactions on Visualization and Computer Graphics 11 (2005), no. 3, 243-253.
[BGO05] Jean-Daniel Boissonnat, Leonidas J. Guibas, and Steve Oudot, Learning smooth objects by probing, Symp. Comp. Geom., 2005, pp. 198-207.
[BO05] Jean-Daniel Boissonnat and Steve Oudot, Provably good sampling and meshing of surfaces, Graphical Models 67 (2005), no. 5, 405-451.
[BO06] , Provably good sampling and meshing of lipschitz surfaces, Symp. Comp. Geom., 2006, pp. 337-346.
[Boo75] William M. Boothby, An introduction to differentiable manifolds and riemannian geometry, Academic Press, New York, 1975.
[Bre42] C. A. Bretschneider, Untersuchung der trigonometrischen Relationen des geradlinigen Viereckes., Archiv der Math. 2 (1842), 225-261.
[BS05] Alexander I. Bobenko and Boris A. Springborn, A discrete Laplace-Beltrami operator for simplicial surfaces, arXiv:math.DG/0503219 v1, 2005.
[BS07] , A discrete Laplace-Beltrami operator for simplicial surfaces, Discrete and Computational Geometry 38 (2007), no. 4, 740-756.
[CB97] H. Chen and J. Bishop, Delaunay triangulation for curved surfaces, Meshing Roundtable, 1997, pp. 115-127.
[CD07] Siu-Wing Cheng and Tamal K. Dey, Delaunay edge flips in dense surface triangulations, 2007, arXiv.org:0712.1959v1.
[CDRR04] Siu-Wing Cheng, Tamal K. Dey, Edgar A. Ramos, and Tathagata Ray, Sampling and meshing a surface with guaranteed topology and geometry, Symp. Comp. Geom., 2004, pp. 280-289.
[CdVM90] Y. Colin de Verdière and A. Marin, Triangulations presque équilatérales des surfaces, J. Differential Geom 32 (1990), no. 1, 199-207.
[CG06] Frédéric Cazals and Joachim Giesen, Delaunay triangulation based surface reconstruction, Effective Computational Geometry for Curves and Surfaces (Jean-Daniel Boissonnat and Monique Teillaud, eds.), Springer, 2006.
[Cha03] Raphaëlle Chaine, A geometric convection approach of 3-D reconstruction, Symp. Geometry Processing, 2003, pp. 218-229.
[Cha06] Isaac Chavel, Riemannian geometry, a modern introduction, 2nd ed., Cambridge, 2006.
[Che89] L. Paul Chew, Constrained Delaunay triangulations, Algorithmica 4 (1989), no. 1, 97-108.
[Che93] , Guaranteed-quality mesh generation for curved surfaces, Symp. Comp. Geom., 1993, pp. 274-280.
[CRS98] P. Cignoni, C. Rocchini, and R. Scopigno, Metro: Measuring error on simplified surfaces, Computer Graphics Forum 17 (1998), no. 2, 167-174.
[dBvKOS98] Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Schwarzkopf, Computational geometry. algorithms and applications, Springer-Verlag, 1998.
[dC76] Manfredo P. do Carmo, Differential geometry of curves and surfaces, PrenticeHall, 1976.
[dC92] , Riemannian geometry, Birkhäuser, 1992.
[Del34] B. Delaunay, Sur la sphère vide, Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk 7 (1934), 793-800.
[Dey07] Tamal Dey, Curve and surface reconstruction; algorithms with mathematical analysis, Cambridge University Press, 2007.
[DFKM08] Daniel Dumitriu, Stefan Funke, Martin Kutz, and Nikola Milosavljević, On the locality of extracting a 2-manifold in $\mathbb{R}^{3}$, Proc. of the 11th Scandinavian Workshop on Algorithm Theory (SWAT), 2008, pp. 270-281.
[DG01] Tamal K. Dey and Joachim Giesen, Detecting undersampling in surface reconstruction, Symp. Comp. Geom., 2001, pp. 257-263.
[DHKL01] Nira Dyn, Kai Hormann, Sun-Jeong Kim, and David Levin, Optimizing 3D triangulations using discrete curvature analysis, Mathematical Methods for Curves and Surfaces, Vanderbilt University, 2001, pp. 135-146.
[DHLM05] M. Desbrun, A. N. Hirani, M. Leok, and J. E. Marsden, Discrete exterior calculus, arXiv:math.DG/0508341, 2005.
[DLJ $\left.{ }^{+} 07\right]$ Junfei Dai, Wei Luo, Miao Jin, Wei Zeng, Ying He, Shing-Tung Yau, and Xianfeng Gu, Geometric accuracy analysis for discrete surface approximation, Comput. Aided Geom. Des. 24 (2007), no. 6, 323-338.
[DLR90] Nira Dyn, David Levin, and Samuel Rippa, Data dependent triangulations for piecewise linear interpolation, IMA Journal of Numerical Analysis 10 (1990), 137-154.
[DLR05] Tamal K. Dey, Gang Li, and Tathagata Ray, Polygonal surface remeshing with Delaunay refinement, Meshing Roundtable, 2005, pp. 343-361.
[DLYG06] Junfei Dai, Wei Luo, Shing-Tung Yau, and Xianfeng Gu, Geometric accuracy analysis for discrete surface approximation, Geometric Modeling and Processing, 2006, pp. 59-72.
[DMA02] Mathieu Desbrun, Mark Meyer, and Pierre Alliez, Intrinsic parameterizations of surface meshes, Proceedings of Eurographics, 2002, pp. 209-218.
[DS89] E. F. D'Azevedo and R. B. Simpson, On optimal interpolation triangle incidences, SIAM J. Sci. Statist. Comput. 10 (1989), no. 6, 1063-1075.
[DZM07a] R. Dyer, H. Zhang, and T. Möller, Delaunay mesh construction, Symp. Geometry Processing, 2007, pp. 271-282.
[DZM07b] , Voronoi-Delaunay duality and Delaunay meshes, Symp. Solid and Physical Modeling, 2007, pp. 415-420.
[DZM08] __, Surface sampling and the intrinsic Voronoi diagram, Computer Graphics Forum (Special Issue of Symposium on Geometry Processing 2008) 27 (2008), no. 5, 1393-1402.
[DZM09] __ Gabriel meshes and Delaunay edge flips, SIAM/ACM Conference on Geometric and Physical Modeling, 2009, pp. 295-300.
[Ede01] H. Edelsbrunner, Geometry and topology for mesh generation, Cambridge, 2001.
[Ede04] , Surface reconstruction by wrapping finite point sets in space, Discrete and Computational Geometry 32 (2004), 231-244.
[ELZ00] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian, Topological persistence and simplification, IEEE Symposium on Foundations of Computer Science, 2000, pp. 454-463.
[EM90] Herbert Edelsbrunner and Ernst Peter Mcke, Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms, ACM Trans. Graph 9 (1990), 66-104.
[EM94] Herbert Edelsbrunner and Ernst P. Mücke, Three-dimensional alpha shapes, ACM Trans. Graph. 13 (1994), no. 1, 43-72.
[Eri01] Jeff Erickson, Nice point sets can have nasty delaunay triangulations, Symp. Comp. Geom., 2001, pp. 96-105.
[ES92] H. Edelsbrunner and N. R. Shah, Incremental topological flipping works for regular triangulations, Symp. Comp. Geom., 1992, pp. 43-52.
[ES94] Herbert Edelsbrunner and Nimish R. Shah, Triangulating topological spaces, Symp. Comp. Geom., 1994, pp. 285-292.
[ES06] Sharif Elcott and Peter Schröder, Building your own DEC at home, ACM SIGGRAPH courses, 2006, pp. 1-5.
[ET93] Herbert Edelsbrunner and Tiow Seng Tan, An upper bound for conforming Delaunay triangulations, Discrete and Computational Geometry 10 (1993), no. 2, 197-213.
[Fed59] Herbert Federer, Curvature measures, Trans. Amer. Math. Soc. 93 (1959), 418-491.
[Flo98] Michael S. Floater, Parametric tilings and scattered data approximation, Int. J. of Shape Modeling 4 (1998), 165-182.
[FSBS06] Matthew Fisher, Boris Springborn, Alexander I. Bobenko, and Peter Schröder, An algorithm for the construction of intrinsic Delaunay triangulations with applications to digital geometry processing, SIGGRAPH Courses, 2006, pp. 6974.
[GGOW08] Jie Gao, Leonidas J. Guibas, Steve Y. Oudot, and Yue Wang, Geodesic Delaunay triangulation and witness complex in the plane, Symp. on Discrete Algorithms, 2008, pp. 571-580.
[GH97] Michael Garland and Paul S. Heckbert, Surface simplification using quadric error metrics, ACM SIGGRAPH, 1997, pp. 209-216.
[Gie99a] Joachim Giesen, Curve reconstruction in arbitrary dimension and the traveling salesman problem, Lecture Notes in Computer Science 1568 (1999), 164-176.
[Gie99b] , Curve reconstruction, the traveling salesman problem and Menger's theorem on length, Symp. Comp. Geom., 1999, pp. 207-216.
[GJ02] Joachim Giesen and Matthias John, Surface reconstruction based on a dynamical system, Computer graphics forum 21 (2002), 363-371.
[GJ03] , The flow complex: a data structure for geometric modeling, Symp. on Discrete Algorithms, 2003, pp. 285-294.
[GKS00] M. Gopi, S. Krishnan, and C. T. Silva, Surface reconstruction based on lower dimensional localized Delaunay triangulation, Computer Graphics Forum (Eurographics) (M. Gross and F. R. A. Hopgood, eds.), vol. 19(3), 2000.
[Gli05] David Glickenstein, Geometric triangulations and discrete Laplacians on manifolds, 2005, arxiv:math.MG/0508188v1.
[GR04] Leonidas Guibas and Daniel Russel, An empirical comparison of techniques for updating Delaunay triangulations, Symp. Comp. Geom., 2004, pp. 170-179.
[GS69] K. Ruben Gabriel and Robert R. Sokal, A new statistical approach to geographic variation analysis, Systematic Zoology 18 (1969), no. 3, 259-278.
[GW03] E. M. Gertz and S. J. Wright, Object-oriented software for quadratic programming, ACM Trans. on Math. Software 29 (2003), 58-81.
[GWM01] S. Gumhold, X. Wang, and R. MacLeod, Feature extraction from point clouds, Meshing Roundtable, 2001, pp. 293-305.
[Hat02] A. Hatcher, Algebraic topology, Cambridge University Press, 2002.
[Hir03] Anil N. Hirani, Discrete exterior calculus, Ph.D. thesis, Caltech, 2003.
[HPW06] Klaus Hildebrandt, Konrad Polthier, and Max Wardetzky, On the convergence of metric and geometric properties of polyhedral surfaces, Geometriae Dedicata 123 (2006), no. 1, 89-112.
[ILTC01] C. Indermitte, Th. Liebling, M. Troyanov, and H. Clémençon, Voronoi diagrams on piecewise flat surfaces and an application to biological growth, Theor. Comput. Sci. 263 (2001), no. 1-2, 268-274.
[KA08] Y. J. Kil and N. Amenta, GPU-assisted surface reconstruction on locallyuniform samples, Meshing Roundtable, 2008, pp. 369-385.
[Kin02] Simon King, Regular flip equivalence of surface triangulations, Topology and its applications 127 (2002), 169-173.
[Kli95] Wilhelm P. A. Klingenberg, Riemannian geometry, second ed., Walter de Gruyter, 1995.
[KR85] D. G. Kirkpatrick and J. D. Radke, A framework for computational morphology, Computational Geometry (G. T. Toussaint, ed.), North-Holland, 1985, pp. 217-248.
[KS98] R. Kimmel and J. A. Sethian, Computing geodesic paths on manifolds, Proceedings of National Academy of Sciences, vol. 95(15), 1998, pp. 8431-8435.
[KSO04] R. Kolluri, J.R. Shewchuk, and J.F. O'Brien, Spectral surface reconstruction from noisy point clouds, Symp. Geometry Processing, 2004, pp. 11-21.
[KW88] Rolf Klein and Derick Wood, Voronoi diagrams based on general metrics in the plane, STACS '88: Proceedings of the 5th Annual Symposium on Theoretical Aspects of Computer Science, Springer-Verlag, 1988, pp. 281-291.
[Lam94] Timothy Lambert, The Delaunay triangulation maximizes the mean inradius, Proc. 6th Canadian Conf. on Computational Geometry, 1994, pp. 201-206.
[Law72] C. L. Lawson, Generation of a triangular grid with application to contour plotting, Section 914 Memo 299, California Instutute of Technolgy Jet Propulsion Laboratory, 1972.
[Law77] , Software for $C^{1}$ surface interpolation, Math. Software III (J. R. Rice, ed.), Academic Press, New York, 1977, pp. 161-194.
[Lei99] Gregory Leibon, Random Delaunay triangulations, the Thurston-Andreev theorem, and metric uniformization, Ph.D. thesis, UCSD, 1999, arXiv:math/0011016v1.
[LL00] Greg Leibon and David Letscher, Delaunay triangulations and Voronoi diagrams for Riemannian manifolds, Symp. Comp. Geom., 2000, pp. 341-349.
[LS03] Francois Labelle and Jonathan Richard Shewchuk, Anisotropic voronoi diagrams and guaranteed-quality anisotropic mesh generation, Symp. Comp. Geom., 2003, pp. 191-200.
[LZ06] J. Li and H. Zhang, Nonobtuse remeshing and decimation, Symp. Geometry Processing, 2006, pp. 235-238.
[MDSB03] Mark Meyer, Mathieu Desbrun, Peter Schröder, and Alan H. Barr, Discrete differential-geometry operators for triangulated 2-manifolds, Visualization and Mathematics III (Hans-Christian Hege and Konrad Polthier, eds.), SpringerVerlag, Heidelberg, 2003, pp. 35-57.
[MR06] Wolfgang Mulzer and Günter Rote, Minimum weight triangulation is NP-hard, Symp. Comp. Geom., 2006, pp. 1-10.
[MT04] Jean-Marie Morvan and Boris Thibert, Approximation of the normal vector field and the area of a smooth surface, Discrete \& Computational Geometry 32 (2004), no. 3, 383-400.
[Mun84] James R. Munkres, Elements of algebraic topology, Addison-Wesley, 1984.
[Mus97] Oleg R. Musin, Properties of the Delaunay triangulation, Symp. Comp. Geom., 1997, pp. 424-426.
[Neg94] S. Negami, Diagonal fips in triangulations of surfaces, Discrete Mathematics 135 (1994), no. 1-3, 225-232.
[OBSC00] Atsuyuki Okabe, Barry Boots, Kokichi Sugihara, and Sung Nok Chiu, Spatial tessellations: Concepts and applications of Voronoi diagrams, second ed., John Wiley, 2000.
[OI03] Kensuke Onishi and Jin-ichi Itoh, Estimation of the necessary number of points in Riemannian Voronoi diagram, Proc. 15th Canad. Conf. Comput. Geom., 2003, pp. 19-24.
[O'R81] Joseph O'Rourke, Polyhedra of minimal area as 3D object models, Intl. Joint Conf. on AI, 1981, pp. 664-666.
[PB01] Sylvain Petitjean and Edmond Boyer, Regular and non-regular point sets: Properties and reconstruction, Computational Geometry: Theory and Applications 19 (2001), no. 2-3, 101-126.
[PC03] Gabriel Peyré and Laurent Cohen, Geodesic remeshing using front propagation, Proceedings IEEE Workshop on Variational and Level Set Methods, 2003, pp. 33-40.
[PC06] , Geodesic remeshing using front propagation, International Journal on Computer Vision 69 (2006), no. 1, 145-156.
[PP93] Ulrich Pinkall and Konrad Polthier, Computing discrete minimal surfaces and their conjugates, Experimental Mathematics 2 (1993), 15-36.
[PSS01] Emil Praun, Wim Sweldens, and Peter Schröder, Consistent mesh parameterizations, ACM SIGGRAPH, 2001, pp. 179-184.
[Raj91] V. T. Rajan, Optimality of the Delaunay triangulation in $\mathbb{R}^{d}$, Symp. Comp. Geom., 1991, pp. 357-363.
[Raj94] , Optimality of the Delaunay triangulation in $\mathbb{R}^{d}$, Discrete and Computational Geometry 12 (1994), 189,202.
[RI97] María-Cecilia Rivara and Patricio Inostroza, Using longest-side bisection techniques for the automatic refinement of Delaunay triangulations, Int. J. Num. Meth. in Eng. 40 (1997), 581-597.
[Rip90] S. Rippa, Minimal roughness property of the Delaunay triangulation, Comput. Aided Geom. Des. 7 (1990), no. 6, 489-497.
[Riv90] Igor Rivin, Euclidean structures on simplicial surfaces and hyperbolic volume, Annals of Mathematics 139 (1990), 553-580.
[RS07] Edgar A. Ramos and Bardia Sadri, Geometric and topological guarantees for the wrap reconstruction algorithm, Symp. on Discrete Algorithms, 2007, pp. 1086-1095.
[Rup93] Jim Ruppert, A new and simple algorithm for quality 2-dimensional mesh generation, SODA '93: Proceedings of the fourth annual ACM-SIAM symposium on discrete algorithms (Philadelphia, PA, USA), Society for Industrial and Applied Mathematics, 1993, pp. 83-92.
[Rup95] J. Ruppert, A Delaunay refinement algorithm for quality 2-dimensional mesh generation, J. of Algorithms 18 (1995), no. 3, 548-585.
[She97] Jonathan Richard Shewchuk, Delaunay refinement mesh generation, Ph.D. thesis, School of Computer Science, Carnegie Mellon University, 1997, Technical Report CMU-CS-97-137.
[She02] , What is a good linear finite element? - interpolation, conditioning, anisotropy, and quality measures, http://www.cs.berkeley.edu/~jrs/papers/elemj.pdf (last viewed 2008), 2002.
[Sib78] R. Sibson, Locally equiangular triangulations, Computer Journal 21 (1978), 243-245.
[Sib81] _ A brief description of natural neighbour interpolation, Interpreting multivariate data (V. Barnet, ed.), John Wiley \& sons, 1981, pp. 21-36.
[Top06] Victor A. Toponogov, Differential geometry of curves and surfaces, Birkhäuser, 2006.
[Tor70] L.V. Toralballa, A geometric theory of surface area, Monatshefte für Mathematik 74 (1970), no. 5, 462-476.
[Tur91] Greg Turk, Generating textures on arbitrary surfaces using reaction-diffusion, ACM SIGGRAPH, 1991, pp. 289-298.
[Tut63] W. T. Tutte, How to draw a graph, Proc. London Math. Society 13 (1963), 743-768.
[vDA95] R. van Damme and L. Alboul, Tight triangulations, Mathematical Methods for Curves and Surfaces, Vanderbilt University Press, 1995, pp. 517-526.
[VHGR08] Evan Vanderzee, Anil N. Hirani, Damrong Guoy, and Edgar Ramos, Wellcentered triangulation, Tech. Report UIUCDCS-R-2008-2936, Department of Computer Science, University of Illinois at Urbana-Champaign, February 2008, Also on arXiv: arXiv:0802.2108v1 [cs.CG].
[VL08] B. Vallet and B. Lévy, Spectral geometry processing with manifold harmonics, Computer Graphics Forum (Special Issue of Eurographics) 27 (2008), no. 2, 251-260.
[Vor07] G. Voronoi, Nouvelles applications des paramètres continus à la théorie des formes quadratiques, J. Reine Angew. Math. 133 (1907), 97-178.
[Vor08] , Nouvelles applications des paramètres continus à la théorie des formes quadratiques, J. Reine Angew. Math. 134 (1908), 198-287.
[WMKG07] Max Wardetzky, Saurabh Mathur, Felix Kälberer, and Eitan Grinspun, Discrete Laplace operators: no free lunch, Symp. Geometry Processing, 2007, pp. 33-37.
[ZvKD10] H. Zhang, O. van Kaick, and R. Dyer, Spectral mesh processing, Computer Graphics Forum (2010), ??-??

## Index

$\alpha$-shape, 41
$\beta$-skeleton, 41
$\theta$-smooth, 29
$k$-Gabriel, 78
2-exposed, 123
3-exposed, 123
4-exposed, 123
accumulation point, 117
affinely independent, 11
allowable region, $6,125,179$
ambient isotopic, 48
ambient space, 16,22
aspect ratio, 17
asymmetric case, 84
bijective, 22
canonical representative, 82
carrier, 12
cell complex, 58
Chavel's lemma, 109
chordal spheroid, 125
circumcentre, 78
circumcentric dual cell, 57
circumcentric dual complex, 3
circumcylinder, 128, 183
circumradius, 78
close vertex, 128, 186
closed ball property, 5, 31
cocone, 38
cocone mesh, 39
cocone triangle, 39
complementary hinge, 83
concentric shells, 116
cone point, 25
conforming homeomorphism, 74
conjugate point, 99
consistent normals, 81
convex, 97
cotan formula, 3, 179
cotan operator, 57
covering, 30
closed, 30
open, 30
cross-sectional diagram, 82
curve on $S, 23$
cut locus, 99
Delaunay complex
intrinsic, 31
manifold, 67 with boundary, 67
restricted, 31
Delaunay mesh, xv
Delaunay refinement, 17
Delaunay tessellation, 59
Delaunay triangle, 12
Delaunay triangulation, 12, 59
conforming, 18, 116
constrained, 18
intrinsic, 2, 35, 59
restricted, iii, 1,33
weighted, 19
density-based sampling criterion, 48
diametric $d$-ball, 78
diffeomorphism, 22
dihedral angle, 81
dimension
simplicial complex, 12
Dirichlet energy, 14
duality structure, 64
edge, 11, 58
boundary, 28
embedded, 115
incident, 125
opposing, 13, 69
original, 115
physical, 115
planar, 116
subtending, 125
unflippable, 5, 93, 114, 123
edge flip, 13
Delaunay, 14
Delaunay extrinsic, 68
edge split, 115
embedding, 22
empty circumdisk, 12
geodesic, 35
immersed, 58
empty diametric ball, 76
equivalent, 82
Euclidean domain, 16
exponential map, 98
exterior half-space, 80,89
extrinsic, 24
face
$k$-face, 11
far vertex, 128, 187
feasible point, 179
feasible region, 125,179
Fenchel-Reshetnyak Theorem, 108
flip-quad, 13
flip-tet, 68
flow complex, 42
Gabriel certificate, 81
Gabriel complex, 20
Gabriel mesh, 4, 41, 79
Gabriel property, 76
Gaussian curvature, 24, 96
Gaussian curvature radius, 96
general intersection property, 32, 168
general position, $9,32,59$
geodesic, 24
closed, 99
periodic, 99
geodesic circumdisk, 2
geodesic disk, 35
geodesic distance, 2, 24
geodesic hinge, 171
geodesic loop, 99
geometric realization, 28
Girard's Theorem, 164
harmonic index, 15
Hausdorff distance, 49
symmetric, 50
hinge, 68
opposing, 69
homeomorphism, 22
horizontal plane, 180
iDt-mesh, 2, 35
immersed empty disk, 58
immersion, 22
incident, 88
injective, 22
injectivity radius, 98
interior, 60
interior half-space, 80,88
interpolation
natural neighbour, 10
nearest neighbour, 10
intrinsic, 1
intrinsic Delaunay triangulation mesh, 35
intrinsic metric, 24
intrinsic sampling radius, $\mathbf{1 0 0}$
isometry, 22
isomorphism, 22
iVd (intrinsic Voronoi diagram), 30
Klingenberg Theorem, 99
lfs (local feature size), 46, 95
lfs $\epsilon$-sample, 48
Lipschitz continuous, 47
local feature size, 46, 95
local reach, 50, 95
local thickness, 39
locally defined surface representation, 77
locally Delaunay, 13, 60
locally Gabriel, 81
loose $\epsilon$-sample, 112
lune, 164
maximal curvature, 24,96
maximal curvature radius, 96
medial axis, 46
discrete, 39
medial ball, 46
meshing, 16
middle plane, 35
Morse-Schönberg Theorem, 107
nerve, 30
nice metrics, 20
nlD (not locally Delaunay), 13
non-sharp, 81, 137, 138
normal convergence, 50
normal curvature, 24
one ring, 27
one skeleton, 58
opposite-side case, 127
orthogonal shadow, 138
piecewise flat surface, 24
pivot, 68
planar contours, 41
pointwise convergence, 50
pole, 38
pole vector, 38
power cell, 19
power diagram, 19, 43
power distance, 19
principle curvatures, 24
properly triangulated, 25
pseudo-disks, 91, 104
pwf surface, 24

Rauch comparison theorem, 171
rDt (restricted Delaunay triangulation), 33
reach, 95
regular interpolant, 39
regular surface, 23
relative curvature, 111
rightness, 53
rVd (restricted Voronoi diagram), 30
same-side case, 127
sample, 9
sampling radius, 48
Schwarz lantern, 45
scr (strong convexity radius), 97
sDm (self-Delaunay mesh), 67
self-Delaunay mesh, iii, xv, 1,67
smooth, 70
with boundary, 67
self-rDt, 75
sharp, 138
simplex, 11, 28
$k$-simplex, 11
degenerate, 11
simplicial complex, 11
abstract, 28
sites, 9
sizing function, 46
sliver tetrahedra, 39, 145
smooth surface, 23
space curve, 23
spherical triangle, 163
split vertex, 116
star, 26, 154
strong convexity radius, 97
strongly convex, 97
surface meshing, 28
surface reconstruction, 28
surjective, 22
symmetric case, 84
tessellation, 58
tetrahedral obstruction, 78
topological persistence, 49
triangle mesh, 26
embedded, 26
manifold, 26
nonobtuse, 44
singular, 26
smooth, 29
well centred, 44
without boundary, 27
triangulates, 29
triangulation, 12, 26, 58
extrinsic, 65
geodesic, 35
minimal weight, 16
proper, 26, 56
regular, 19
regular (of pwf surface), 43
tangled, 148
weighted (of pwf surface), 43
tubular neighbourhood, 50, 95
umbrella, 27
full, 28
smooth, 29
vertex, 11, 58
boundary, 27
original, 115
smooth, 29
split, 115
vertical, 180
Voronoi cell, 8, 60
intrinsic, 30
restricted, 30
Voronoi diagram, 9
intrinsic, 30, 60
restricted, 5, 30
well formed, 33
Voronoi edge, 9, 60
associated, 63
internal, 60
Voronoi face, 9
Voronoi facet, 9
Voronoi neighbour, 9

Voronoi vertex, 9, 60
atomic, 168
Willmore energy, 53


[^0]:    ${ }^{1}$ Image from Aliez et al. "Variational Tetrahedral Meshing" SIGGRAPH 2005, reproduced with kind permission of the authors.

[^1]:    ${ }^{2}$ Reprinted from http://dx.doi.org/10.1016/S0925-7721(01)00047-5 Computational Geometry, Volume 22, Number 1-3, Jonathan Richard Shewchuk, Delaunay refinement algorithms for triangular mesh generation, pages 21-74, 2002, with permission from Elsevier.

[^2]:    ${ }^{3}$ Although we have not explicitly defined a pwf surface with boundary, this is not an issue here because the definition of a triangle mesh only involves a triangulated pwf surface, i.e., a manifold simplicial complex.
    ${ }^{4}$ This paper is also the one in which the phrase "digital geometry processing" was coined.

[^3]:    ${ }^{5}$ An alternative to geodesic curves is mentioned in [LL00]. They propose to employ "middle planes" (a middle plane between $a$ and $b$ on $S$ is given by $\left\{x \in S \mid d_{S}(x, a)^{2}-d_{S}(x, b)^{2}=c^{2}\right\}$ for some constant $c$ ). This approach has appeal in that it is amenable to extensions to higher dimensions, however the details of their proposal were never published.

[^4]:    ${ }^{6}$ These figures are inspired by figures in [KSO04] and [ABK98] respectively.

[^5]:    ${ }^{7}$ In earlier work on meshing Euclidean domains, Ruppert [Rup95] introduced a different quantity which he called the local feature size. Ruppert's definition of local feature size is still widely employed in the meshing community, whereas Amenta and Bern's definition is the norm in geometry processing. We make no attempt to remedy this unfortunate situation here. In this thesis, "local feature size" (lfs) always refers to the sizing function defined by Amenta and Bern [AB98].

[^6]:    ${ }^{1}$ Some authors define the dihedral angle to be the angle between the faces themselves - the supplement of our dihedral angle. Our definition is convenient and not unprecedented. See [CD07], for example.

[^7]:    ${ }^{1}$ Thanks to A. Bobenko for this insight.

[^8]:    ${ }^{1}$ Thanks to Karan Singh for pointing me towards these results.

[^9]:    ${ }^{1}$ Thanks to Jonathan Shewchuk for bringing the infinite genus case to my attention.

