

ABSTRACT

Title of Dissertation: COMPUTATIONAL GEOMETRY IN THE
FUNK AND HILBERT METRICS

Auguste Henry Gezalyan
Doctor of Philosophy, 2026

Dissertation Directed by: Professor David Mount
Department of Computer Science

The Hilbert metric is a generalization of the Cayley-Klein model of hyperbolic geometry to arbitrary convex bodies. It has applications in convex approximation, quantum information theory, machine learning, and real analysis. The Hilbert metric can be constructed as the average of the forward Funk and reverse Funk metrics, which are asymmetric distance functions with their own applications in convex approximation and real analysis. Despite the significance of these metrics, there has been little work on understanding the Hilbert and Funk metrics from a computational geometry perspective. This dissertation fills this gap by examining their fundamental geometric structure and provides efficient algorithms for geometric problems in these metrics.

In the Hilbert metric, we provide several contributions on 2-dimensional polygons with m sides and n sites. We characterize bisectors as piecewise conics and provide an $O(mn \log n)$ expected time algorithm for constructing Voronoi diagrams in the metric, matching the worst-case $\Theta(mn)$ complexity up to a logarithmic factor. We analyze the behavior of these diagrams in their

dynamic setting and study their discontinuities. We present an $O(n(\log n + \log^3 m))$ expected time algorithm for Delaunay triangulations in the Hilbert metric. Unlike the Euclidean case, the Hilbert Delaunay triangulation does not cover the convex hull of the point set. We introduce the concept of the Hilbert hull to characterize what is actually covered and provide an $O(nh \log^2 m)$ time algorithm for computing it, where h is the number of boundary points on the hull.

We extend our analysis to the forward and reverse Funk metrics in d -dimensional elliptical cones and 3-dimensional polygonal cones. We prove that bisectors in these metrics consist of rays from the cone's apex and show a reduction proving the equivalence between d -dimensional Funk Voronoi diagrams and $(d - 1)$ -dimensional additively weighted Funk Voronoi diagrams. Leveraging this, we present an $O(n^{\lceil (d-1)/2 \rceil + 1})$ time algorithm for d -dimensional elliptical cones and an $O(mn \log n)$ time algorithm for 3-dimensional polygonal cones. We also provide a complete characterization of when three sites admit circumcenters in 3-dimensional polygonal cones.

COMPUTATIONAL GEOMETRY IN THE FUNK AND HILBERT METRICS

by

Auguste Henry Gezalyan

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2026

Advisory Committee:

Professor David M. Mount, Chair/Advisor
Professor William Gasarch
Professor Michelle Mazurek
Professor Erin Molloy
Professor David Weintrop

© Copyright by
Auguste Henry Gezalyan
2026

Dedication

Dedicated to Tristan Alexandre Lascoumes.

Acknowledgments

I would like to acknowledge the work done by the students I've worked with over the semesters at UMD on Hilbert Geometry and other projects: Madeline Bumpus, Caesar Dai, Sam Munoz, Renita Santhoshkumar, Songyu Ye, Soo Kim, Carlos Lopez, Daniel Skora, Zofia Stefankovic, Hridhaan Banerjee, Carmen Isabel Day, Megan Hunleth, Sarah Hwang, Olya Golovatskaia, Nithin Parepally, Lucy Wang, Ainesh Chatterjee, Hongyang Du, Sukrit Mangla, Kenny Wu, Khawaja Khalid, Samuel Kushnir, Julian Vanecek, Soren Brown, June Cagan, Veena Kailad, Chaewoon Kyoung, Rowan Shigeno, Yasmine Tajeddin, and Andrew Waggoner. I would also like to thank my excellent advisor Professor David Mount for his mentorship, Professor Bill Gasarch for the REU-CAAR opportunities, and Professors Erin Molloy, Michelle Mazurek, and David Weintraub, for advice and being on my committee.

Table of Contents

Dedication	ii
Acknowledgments	iii
Table of Contents	iv
List of Figures	vii
Chapter 1: Introduction	1
1.1 Overview	1
1.2 Preliminaries	4
1.3 Voronoi Diagrams in the Hilbert Metric	5
1.4 Analysis of Dynamic Voronoi Diagrams in the Hilbert Metric	5
1.5 Delaunay Triangulations in the Hilbert Metric	6
1.6 Voronoi Diagrams in the Funk Metric	7
Chapter 2: Preliminaries	9
2.1 Metric Spaces	9
2.2 Hyperbolic Geometry	10
2.3 Funk and Hilbert Metrics	11
2.4 Voronoi diagrams and Delaunay Triangulations	14
2.5 Conical Funk	16
Chapter 3: Voronoi Diagrams in the Hilbert Metric	21
3.1 Introduction	21
3.2 Hilbert Metric Balls in Higher Dimensions	22
3.2.1 Hilbert Balls of Points	22
3.2.2 Hilbert Balls of Lines and Line Segments	27
3.3 Characterizing Voronoi Diagrams in the Hilbert Metric	31
3.4 Bisectors in the Hilbert Metric	32
3.4.1 Bisector Segments and Combinatorial Complexity	34
3.4.2 Bisector Segments Structure	36
3.5 Randomized Incremental Algorithm	40
3.5.1 The Hilbert Metric and Abstract Voronoi Diagrams	40
3.5.2 Point Sites	40
3.5.3 Segment Sites	44

3.6	Divide and Conquer Algorithm	45
Chapter 4:	Analysis of Dynamic Voronoi Diagrams in the Hilbert Metric	51
4.1	Introduction	51
4.2	Hilbert Bisector analysis	51
4.2.1	Finding Bisectors in a Sector	54
4.3	Analysis of Hilbert Voronoi Diagrams	61
Chapter 5:	Delaunay Triangulations in the Hilbert Metric	66
5.1	Introduction	66
5.2	The Hilbert Delaunay Triangulation	67
5.3	Hilbert Bisectors	70
5.4	Hilbert Circumcircles	75
5.4.1	Balls at infinity	75
5.5	Computing Circumcircles	78
5.6	Building the Triangulation	81
5.6.1	Orienting and Augmenting the Triangulation	82
5.6.2	Local and Global Delaunay	84
5.6.3	Incremental Construction	85
5.7	The Hilbert Hull	92
Chapter 6:	Voronoi Diagrams in the Funk Metric	94
6.1	Introduction	94
6.2	Computing Funk Voronoi Diagrams in Cones	95
6.2.1	Empty Voronoi Cells	95
6.2.2	Voronoi Cells	96
6.2.3	Voronoi Algorithms	98
6.2.4	Voronoi Diagrams in d-Dimensional Elliptical Cones	102
6.2.5	Voronoi Diagrams in 3-Dimensional Polygonal Cones.	106
6.2.6	Circumcenter Characterization in 3-Dimensional Cones	110
Chapter 7:	Conclusions	115
	Bibliography	118

List of Figures

2.1	(a) The Funk weak metric and (b) the Hilbert metric.	12
2.2	(a) The Voronoi diagram of a set of points and (b) its Delaunay triangulation. . . .	14
2.3	(a) An elliptical cone. (b) A polygonal cone.	16
2.4	(a) A collection of 5 points in a cone C . (b) The cone C translated to the points. (c) A directed Hasse diagram describing their partial order, $d \leq_C e, c \leq_C e, b \leq_C$ $c, a \leq_C b$	17
2.5	Two points a and b where (a) $F_C(a, b) < 0$, (b) $F_C(a, b) = 0$, (c) $F_C(a, b) > 0$	18
3.1	The Hilbert ball of a point p	23
3.2	The Hilbert ball of a point p on $\partial \Omega$ through a point q	26
3.3	The closest point to q on line ℓ	28
3.4	The Hilbert ball of a line ℓ	30
3.5	Three-step process for building the Hilbert ball of a line segment $\overline{ab}$	31
3.6	Conditions for a point z to lie on the Hilbert bisector between sites p and p'	33
3.7	(a) The sectors of p with respect to Ω and (b) the Hilbert bisector (in red) between p and p'	33
3.8	Proof of Lemma 3.4.4.	36
3.9	Segments of a point-point bisector	37
3.10	(a) Closest point on b to point q , (b) point-segment bisector, (c) segment-segment bisector	39
3.11	Inserting a new site in the randomized algorithm.	41
3.12	Divide-and-conquer algorithm, bisectors, and the proof of Lemma 3.6.1(iii).	46
3.13	Tracing $B(S_L, S_R)$. (Drawing is not geometrically accurate.)	48
4.1	Illustration of different <i>sectors</i>	52
4.2	A projection between two quadrilaterals.	55
4.3	Illustration for the proof of Lemma 4.2.6 (a bisector in the three-edge case with shared edge in front of both sites).	58
4.4	The bisector conics for the cases given in Table 4.1 visualized.	60
4.5	The point p is equidistant from both sites s and t by the invariance of the cross ratio. . . .	62
4.6	A discontinuity in the Hilbert metric Voronoi.	63
4.7	Illustration of region Z	65
5.1	(a) A convex polygon Ω and sites, (b) the Voronoi diagram, and (c) the Delaunay triangulation.	67
5.2	Proof of Lemma 5.2.2.	69
5.3	Properties of points on the Hilbert bisector.	71

5.4	Computing where a ray through a site hits the bisector.	72
5.5	Computing the endpoints of the (p, q) -bisector.	73
5.6	Constructing the Hilbert ball $B_u(x, p)$ at infinity.	76
5.7	Proof of Lemma 5.4.2.	77
5.8	The overlap region $Z(p, q)$ and the outer region $W(p, q)$	78
5.9	Circumcircles and bisector intersection.	79
5.10	Computing the center of $\triangle pqr$	80
5.11	The augmented triangulation.	82
5.12	Proof of Lemma 5.6.2.	85
5.13	(a) Initialization and insertion into a (b) standard triangle, (c) tooth, (d) gap.	86
5.14	(a) Standard-triangle edge flip, (b) tooth edge flip, and (c) fixing a tooth.	88
5.15	Proof of Lemma 5.6.3.	89
6.1	(a) Voronoi cells in the reverse Funk metric, and (b) in the forward Funk metric.	97
6.2	(a) Reverse Funk ball tangent to two forward Funk balls, (b) forward Funk ball tangent to two reverse Funk balls.	103
6.3	(a) The apex of C' lies in C . (b) Locating the reverse Funk vertex v in B	104
6.4	(a) Spokes in a forward Funk Voronoi and, (b) reverse Funk Voronoi.	107
6.5	In the reverse Funk, (a) Cross section of a circular cone, (b) (c) (d) regions labeled in a circular cone with differently weighted sites, (e) regions labeled in a polygonal cone, (f) pencil of balls along the bisector in a polygonal cone.	113

Chapter 1: Introduction

1.1 Overview

The Hilbert metric is a generalization of the Cayley-Klein model of hyperbolic geometry to arbitrary convex bodies. It was introduced by David Hilbert in 1895 [1]. It has applications in convex approximation, quantum information theory, machine learning, and real analysis among other fields. Given any convex body Ω in d -dimensional space, the Hilbert metric defines a distance between any pair of points in the interior of Ω (see Chapter 2 for definitions). The Hilbert metric has a number of useful properties. It is invariant under projective transformations, and straight lines are geodesics. When Ω is a Euclidean ball, it realizes the Cayley-Klein model of hyperbolic geometry, and when Ω is a simplex, it provides a natural metric over discrete probability distributions (see, e.g., Nielsen and Sun [2, 3]). An excellent resource on Hilbert geometries is the “Handbook of Hilbert Geometry” by Papadopoulos and Troyanov [4].

The Hilbert geometry provides new insights into classical questions from convexity theory. Efficient approximation of convex bodies has a wide range of applications, including approximate nearest neighbor searching both in Euclidean space [5] and more general metrics [6], optimal construction of ε -kernels [7], solving the closest vector problem approximately [8–11], and computing approximating polytopes with low combinatorial complexity [12, 13]. These works all share one thing in common—they approximate a convex body by covering it with elements that behave much

like metric balls. These covering elements go under various names: Macbeath regions, Macbeath ellipsoids, Dikin ellipsoids, and $(2, \varepsilon)$ -covers. Vernicos and Walsh showed that these shapes are, up to a constant scaling factor, equivalent to Hilbert balls [14–16]. In addition, the Hilbert metric behaves nicely in the context of flag approximability of convex polytopes, as studied by Vernicos and Walsh [17].

The Hilbert geometry has proven to be of use in the field of networking. It has been shown that complex networks have an underlying hyperbolic geometry [18–20]. Recently, F. Nielsen and K. Sun studied non-linear embeddings in the Hilbert simplex geometry [3] and found them to be competitive. Other applications of the Hilbert metric include machine learning [2, 3], quantum information theory [21], real analysis [22], and optimal mass transport [23].

Many of these applications are due to the Hilbert metric being a natural metric in the simplex geometry and on convex cones. A simplex is the d -dimensional generalization of a triangle. Many problems can be modeled geometrically on a simplex, for instance probability distributions over a discrete set of $d + 1$ outcomes can be modeled as points in the d -dimensional simplex. The Hilbert metric forms a natural distance in this context isometric to the induced norm on a hexagonal grid. It has been shown to be competitive with Kullback–Leibler divergence for clustering applications [2]. In convex cones the Hilbert metric provides a natural distance between rays, similar to the Thompson metric. Details on its applications related to convex cones and analysis can be found in the survey by Lemmens and Nussbaum [22].

Despite its obvious appeals, only recently has there been any work on developing classical computational geometry algorithms that operate in the Hilbert metric. Nielsen and Shao characterized balls in the Hilbert metric defined by a convex polygon with m sides [24]. They showed that Hilbert balls are convex polygons bounded by $2m$ sides, and they can be computed in $O(m)$ time.

They also developed dynamic software for generating them. At this point, it should be noted that the Hilbert metric can be constructed as the average of the forward Funk and reverse Funk metrics, which are asymmetric distance functions with their own applications in convex approximation and real analysis. In contrast to the Hilbert metric, the Funk metrics are completely unstudied in computational geometry.

The Funk metrics were originally defined by Paul Funk in 1929 to study geometries with straight line geodesics [25] (see Chapter 2 for definitions). This work directly addresses Hilbert’s fourth problem, which sought to characterize all geometries where straight lines are geodesics. The forward Funk and reverse Funk metrics are deeply connected to the Hilbert metric. As previously stated, the Hilbert metric can be defined as the average of the forward Funk and reverse Funk metrics. The Funk metrics also play a role in the definition of other distance functions, such as the Thompson metric [26], which is defined as the maximum of the forward Funk and reverse Funk metrics.

As the Hilbert geometry has many applications, this motivates the study of the Funk metrics in two ways. First, understanding the behavior of the Funk geometry provides insights for constructing improved algorithms in related metrics, such as Hilbert and Thompson. Second, there is evidence that the Funk metrics are more natural to study than Hilbert. Faifman argues this perspective in “A Funk perspective on billiards, projective geometry and Mahler volume” [27]. Faifman shows that the Holmes-Thompson volume and surface area are preserved under polarity in the Funk metrics with a myriad of other desirable properties.

The Funk metrics have also recently emerged as an area of contemporary research in their own right in Finsler geometries. Funk is considered one of the most basic Finsler geometries [4]. In particular, Papadopoulos and Troyanov showed that every convex body contains a canonical weak

Finsler structure whose distance is inherently the forward Funk metric [28]. Additionally, Shen used it to construct an R-flat spray and employed the forward Funk metric to find a Finsler metric with zero curvature [29]. Most recently, the forward Funk metric was used to solve problems involving trajectory in the context of Zermelo’s navigation problem [30].

Despite the significance of the Hilbert and Funk metrics, as mentioned, there has previously been little work on understanding these metrics from a computational geometry perspective. This dissertation fills this gap by examining the fundamental geometric structure of the Hilbert and Funk metrics and provides efficient algorithms for geometric problems in them. In these next sections, we present our contributions.

1.2 Preliminaries

In Chapter 2, we introduce our preliminaries. We begin by defining metric spaces and balls in metric spaces. After a section describing the hyperbolic metric, we define our three metrics, Funk, reverse Funk, and Hilbert metrics. Next, we introduce the concepts of Voronoi diagrams, bisectors, and Delaunay triangulations (the graphical dual of the Voronoi).

Next, we introduce the conical Funk metrics. The conical Funk metrics are a generalization of the Funk metrics defined earlier to unbounded convex cones. We start by defining a conical partial order on C . Then, we give the conical definition for forward Funk and reverse Funk with respect to that order. Next, we provide an alternative characterization of the conical Funk distances with respect to distances to the boundary of the cone, and we characterize balls in the two conical Funk metrics as translates of the boundary cone. Additionally, we prove that straight lines in the Funk metrics are geodesics.

1.3 Voronoi Diagrams in the Hilbert Metric

In Chapter 3, we begin with a generalization of the characterization of a Hilbert ball to arbitrary dimensions. Given a convex polytope Ω bounded by m facets in $\mathbb{R}^d$, we show that the ball of radius r around a point p in Ω is a convex polytope bounded by m^2 facets. Similarly, we show that in $\mathbb{R}^2$ the Hilbert ball around lines and line segments are convex polygons with complexity at most $2m$ and can be constructed in time $O(m)$.

From there we proceed to Voronoi diagrams in the Hilbert metric. We begin by showing that Voronoi cells in the Hilbert metric are stars (Lemma 3.3.1). Consequently, we are able to bound the complexity of the Voronoi diagrams in the Hilbert metric as $O(mn)$ where n is the amount of sites and m is the amount of edges of Ω . Additionally, we show bisectors are piecewise conic curves. From there, we present three algorithms. The first two algorithms are randomized incremental insertion algorithms. The first of these algorithms is for point sites and the second is for point and line segment sites. Both algorithms run in expected time $O(mn \log(n))$. Lastly, we give a deterministic divide and conquer algorithm for point sites that runs in time $O(mn \log(n))$. These results appeared in the paper “Voronoi Diagrams in the Hilbert Metric” [31].

1.4 Analysis of Dynamic Voronoi Diagrams in the Hilbert Metric

In Chapter 4, we begin by showing how, if given two sites, we can break down a polygon into sectors in which the bisector between the two has a different conic form with respect to the Hilbert metric. We give an explicit equation for the bisector depending on the sector. Then, we analyze the different types of sectors in a case by case fashion. We show that when a sector is

defined by four edges of Ω the bisector is a hyperbola. When the sector is defined by three edges of Ω it can be a straight line or any conic depending on the arrangement of the edges with respect to the two sites. And when the sector is defined by two edges of Ω the bisector is always a straight line.

We then proceed to do two things. First, we identify exactly when a Hilbert bisector degenerates to having a two dimensional component. This happens exactly when two sites lie on a line through a vanishing point of two edges of Ω . Then, we identify a region between two points such that every ball containing those two points on its boundary contains that region, and we show that it has positive measure, thus showing that it is not a degeneracy. This chapter covers the results from the paper “Analysis of Dynamic Voronoi Diagrams in the Hilbert Metric” [32].

1.5 Delaunay Triangulations in the Hilbert Metric

Chapter 5, begins by considering the idea of the Hilbert Delaunay triangulation, which we define to be the dual of the Hilbert Voronoi diagram. We begin by showing that three points in the Hilbert metric define at most one circumcircle, and that the Hilbert Delaunay triangulation of a set of sites spans those sites. We then proceed to prove some supporting lemmas that allow us to perform important subroutines, such as intersecting rays from sites with bisectors, and computing the endpoints of bisectors.

Next, we introduce the idea of balls at infinity through points. These are limit balls whose centers lie along rays or bisectors that intersect the boundary of Ω and pass through a point or pair of points. We use these to augment the Hilbert Delaunay triangulation, which allows us to handle the fact that the boundary of the triangulation is not necessarily convex. Using these infinite balls,

we fully characterize when three balls admit a Hilbert circumcircle.

Using the tools developed in this chapter, we are able to calculate whether three points lie on a Hilbert circumsphere or not in $O(\log^3(m))$ time. This allows us to compute the Hilbert Delaunay triangulation, which we show has complexity $O(n)$. Our algorithm is incremental and runs in randomized expected time $O(n(\log n + \log^3 m))$. We prove the fact that our algorithm is correct using the fact that locally Delaunay and globally Delaunay are the same in the Hilbert metric. We end by showing how to construct the Hilbert hull (the exterior Hilbert Voronoi cells) in time $O(nh \log^2 m)$, where h is the amount of cells on the boundary. These results appeared in the work “Delaunay Triangulations in the Hilbert Metric” [33]

1.6 Voronoi Diagrams in the Funk Metric

In Chapter 6, we move on to the Funk Voronoi diagram. We consider both the forward Funk and reverse Funk weak metrics in both elliptical and polygonal cones. We begin by discussing when a site has or does not have a Voronoi cell in the Funk metrics, and we identify this criterion with respect to the partial order of the sites. We then prove that Voronoi cells in the Funk metrics are stars. In order to deal with d -dimensional Voronoi cells we prove that bisectors in the Funk metrics are composed of rays through the origin and that sites are equivalent on rays up to additive factors. Hence their Voronoi diagrams can be computed on cross sections of the cone. Next, we examine the Funk Voronoi in the d -dimensional case on elliptical cones.

We show that the Voronoi diagram of a set of sites on d -dimensional elliptical cones is closely related to the Apollonius diagram. We show that Voronoi vertices in the Funk correspond to the existence of Voronoi vertices in the Apollonius. We show an equivalent result for reverse Funk,

but only if the Apollonius Voronoi vertex is in the cone. We then provide a transformation and algorithm to Apollonius for the two metrics. We show that the Voronoi diagram of a set of sites in an elliptical cone in the Funk metrics can be computed in time $O(n \log h)$ when $d - 1 = 2$, and $O(n^{\lceil \frac{d-1}{2} \rceil + 1})$ otherwise, where h is the amount of sites with cells, n is the amount of sites, and d is the dimension.

Next, we proceed to examine the Funk Voronoi diagrams in polygonal cones. First, we show that we can divide the cones into sectors and that the bisector between sites depends on the sector they are in. Here, we reduce the abstract Voronoi diagrams and show that Voronoi diagrams in 3-dimensional polygonal cones can be computed in time $O(nm \log(n))$ where n is the amount of sites, and m is the complexity of the polygonal cone.

Lastly, we characterize when three points lie on the boundary of a Ball in the Funk metrics. This turns out to be similar to the criterion in the Hilbert metric, where we had three regions that determined the criterion. In the Funk metrics, since some triples of sites can have two circumcenters, we have five.

Chapter 2: Preliminaries

2.1 Metric Spaces

A *metric space* $(\mathbb{X}, d)$ is an abstraction of the concept of a distance function d on a set of points in a domain $\mathbb{X}$. In order for a function to be considered a distance function for a metric space, it must satisfy three properties. Formally:

Definition 2.1.1 (Metric Space). Given a set $\mathbb{X}$ and a function $d : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}^{\geq 0}$, the pair $(\mathbb{X}, d)$ is said to be a metric space when d satisfies the following properties for arbitrary $x, y, z \in \mathbb{X}$:

1. $d(x, y) = 0 \iff x = y$
2. $d(x, y) = d(y, x)$
3. $d(x, y) + d(y, z) \leq d(x, z)$.

The first property is known as the *identity of indiscernibles*, it says that if two things are zero distance from each other they cannot be distinct in $\mathbb{X}$. The second property, called *symmetry*, states that, given two points x and y the distance from x to y should be the same as from y to x . Distance functions without this property arise frequently and they form what are called *weak metric spaces*. The last property is the *triangle inequality*, which states that the distance from x to z yields the length of the shortest path from x to z . A standard construct in metric spaces is the idea of an *ball* or *neighborhood* of a point in that space. This is often denoted as: $B(x, r)$, the ball of radius r

around a point x .

Definition 2.1.2 (Ball). Given a metric space $(\mathbb{X}, d)$, $B(x, r)$ is said to be the ball of radius r around a point x :

$$B_{\mathbb{X}}(x, r) = \{y \in \mathbb{X} \mid d(x, y) \leq r\}$$

Additionally let $B_{\mathbb{X}}(x : y : z)$ be the ball circumscribing x, y, z . The sub scripted space is omitted when it is clear from context

Metric spaces are a fundamental mathematical tool for analysis in metric or topological spaces. The most common metric space used is the Euclidean metric, which represents the distance between two points in d dimensional Euclidean space. For the rest of the dissertation we will use $d(\cdot, \cdot)$ to refer to the Euclidean distance unless otherwise specified. In computer science, there are many metric spaces that see common use. For instance, the Hamming space with the Hamming distance, the disk with the Cayley-Klein metric, $\mathbb{R}^d$ with any Minkowski distance of order p , the Hausdorff distance on sets in metric spaces, and Frechet space all are metric spaces. In the rest of the dissertation, we will only give the space the metric is taking place in, if it is not obvious, and we abbreviate the Hilbert/Funk/hyperbolic metric space to Hilbert/Funk/hyperbolic metric. When we say Funk metric without the forward or reverse predicate we mean the forward Funk.

2.2 Hyperbolic Geometry

There are three types of Riemannian manifolds of constant curvature, ones with zero curvature (Euclidean), ones with positive curvature (spherical), and ones with negative curvature (hyperbolic). All three of these form metric spaces with their corresponding distances that are naturally present in the world we live in. This includes the hyperbolic metric which is present in Minkowski

spacetime and anti-de Sitter space [18]. Each of these metric spaces has its own unique features that make them useful for modeling in different contexts. Some of the particular characteristics of hyperbolic geometry that separate it from spherical and Euclidean are that balls in hyperbolic space expand exponentially and that given a line l and point p (not on l) in hyperbolic space there are an infinite amount lines passing through p that are parallel to l .

Since models of hyperbolic geometry cannot be isometrically embedded into Euclidean space due to their growth, there are many equivalent models of hyperbolic space that are studied. This includes the Poincaré disk and half-plane model, the hyperboloid model and the Cayley-Klein model. Each of these models comes with unique properties that make them useful. For instance the Poincare disk model is angle preserving but does not preserve lines. The Cayley-Klein model of hyperbolic space on the other hand does preserve lines. The Cayley-Klein model of hyperbolic space models points in hyperbolic space as points in the interior of a Euclidean unit ball paired with a distance metric. It is exactly the Hilbert metric but restricted to the unit ball (Definition 2.3.3). It is important to note that in general, the Hilbert metric is not always Riemannian [4]

2.3 Funk and Hilbert Metrics

Throughout, let Ω denote a *convex body* in $\mathbb{R}^d$, that is, a closed, bounded, full-dimensional convex set. Let $\partial\Omega$ and $\text{int}(\Omega)$ denote its boundary and interior, respectively. Given points $p, q \in \mathbb{R}^d$, let $d(p, q)$ denote the Euclidean distance between these points. Given two distinct points $p, q \in \Omega$, let $\chi(p, q)$ denote the *chord* defined as the intersection of the line passing through p and q with Ω .

Before defining the Hilbert metric, it will be convenient to define a simpler (asymmetric)

distance function called the *Funk metric*, also known as the *Forward Funk metric*.

Definition 2.3.1 (Forward Funk metric (bounded convex bodies)). Given a convex body Ω in $\mathbb{R}^d$ and two distinct points $p, q \in \text{int}(\Omega)$, let y denote the point where a ray shot from p to q intersects $\partial\Omega$. Define the *Funk weak metric* to be

$$F_{\Omega}(p, q) = \ln \frac{d(p, y)}{d(q, y)},$$

and define $F_{\Omega}(p, p) = 0$ (see Figure 2.1(a)).

Definition 2.3.2 (Reverse Funk metric (bounded convex bodies)). Given a convex body Ω in $\mathbb{R}^d$ and two distinct points $p, q \in \text{int}(\Omega)$, define the reverse *Funk weak metric* to be

$$F_{\Omega}^r(p, q) = F_{\Omega}(q, p),$$

and define $F_{\Omega}^r(p, p) = 0$.




Figure 2.1: (a) The Funk weak metric and (b) the Hilbert metric.

Observe that in the limit as q approaches y along the chord $\chi(p, y)$ the Funk distance increases to $+\infty$, and thus, the boundary is infinitely far away from any interior point of Ω . The Funk weak metric is not symmetric, but it satisfies the triangle inequality. Symmetrizing this yields a true metric, called the Hilbert metric [1].

Definition 2.3.3 (Hilbert metric). Given a convex body Ω in $\mathbb{R}^d$ and two distinct points $p, q \in \text{int}(\Omega)$, let x and y denote endpoints of the chord $\chi(p, q)$, so that the points are in the order $\langle x, p, q, y \rangle$. Define the *Hilbert metric* to be

$$H_{\Omega}(p, q) = \frac{F_{\Omega}(p, q) + F_{\Omega}(q, p)}{2} = \frac{1}{2} \ln \frac{d(p, y)d(q, x)}{d(q, y)d(p, x)},$$

and define $H_{\Omega}(p, p) = 0$ (see Figure 2.1(b)).

Hilbert showed that line segments are geodesics, and if q lies on the line segment between p and r , then $H_{\Omega}(p, q) + H_{\Omega}(q, r) = H_{\Omega}(p, r)$. Note however that generally there may be multiple shortest paths between two points, and hence geodesics need not be line segments (see, e.g., [34]). As in the Funk weak metric, the boundary of Ω is infinitely far away from any interior point. Observe that the quantity in the $\ln$ term in the definition is the cross ratio of $(p, q; y, x)$. It follows that the Hilbert metric is invariant under projective transformations.

It is well known that $H_{\Omega}(a, b) + H_{\Omega}(b, c) = H_{\Omega}(a, c)$ if and only if $r_{\Omega}(a, b), r_{\Omega}(b, c), r_{\Omega}(a, c)$ are collinear and $r_{\Omega}(b, a), r_{\Omega}(c, b), r_{\Omega}(c, a)$ are collinear, where $r_{\Omega}(a, b)$ is the point where the ray ab meets $\partial(\Omega)$ [35, Theorem 12.4]. This follows from the Hilbert metric's characterization as being the average of the forward Funk and reverse Funk metrics. It is also well known that geodesics are unique if and only if Ω is strictly convex [35, Corollary 12.7]. For further information, see the first chapters of the handbook on Hilbert geometry by Papadopoulos and Troyanov [4].

2.4 Voronoi diagrams and Delaunay Triangulations

A *Voronoi diagram* of a set of points S , denoted $V(S)$, is a fundamental construction in geometry with a multitude of applications [36]. Given a set of points S in a metric space $(\mathbb{X}, d)$, we want to divide $\mathbb{X}$ into a cellular structure around the points in S , such that each cell represents the set of points closest to S (see 2.2(a)). We call these Voronoi cells.

Definition 2.4.1 (Voronoi Cell). Given a set of sites S , the Voronoi cell of a site $s \in S$ is:

$$V(s) = \{q \in \mathbb{X} \mid d(q, s) \leq d(q, t), \forall t \in S \setminus \{s\}\}.$$

The boundary of these Voronoi cells are formed by bisectors. The bisector between two sites, x and y is the set of points in the metric space equidistant to both of the sites (see 2.2(a)).

Definition 2.4.2 (Bisector). Given two sites x and y , the bisector between the two sites is:

$$(x, y)\text{-bisector} = \{z \in \mathbb{X} \mid d(x, z) = d(y, z)\}.$$




Figure 2.2: (a) The Voronoi diagram of a set of points and (b) its Delaunay triangulation.

The union of the Voronoi cells of the sites forms the Voronoi diagram.

Definition 2.4.3 (Voronoi Diagram). Given a set of sites S , the Voronoi diagram of the set S , $V(S)$, is the cellular decomposition of space into the Voronoi cells of the sites in S .

Note that sites can be given **weights** in a weighted Voronoi diagram, $S = \{(s_1, w_1), \dots, (s_n, w_n)\}$.

In that case, the Voronoi cell of a site s and it's bisector are:

$$V((s, w)) = \{q \in C \mid d(s, q) + w \leq d(s', q) + w' \text{ for all } (s', w') \in S\},$$

$$((s, w), (s', w'))\text{-bisector} = \{z \in \mathbb{X} \mid d(s, z) + w = d(s', z) + w'\}.$$

One of the primary uses of Voronoi diagrams is for nearest neighbor querying. Given a set of sites S the Voronoi diagram of the sites can be used to preprocess the sites into a data structure that allows for, given an arbitrary point, finding the closest site to it in logarithmic time. Their applications come in many different forms, from file searching to crystallography. Aurenhammer [36] provides an excellent resource on Voronoi diagrams and related structures. There are a multitude algorithms to implement Euclidean Voronoi diagrams in a variety of contexts from randomized incremental to plane sweep. A good survey is given by Fortune [37].

The graph-theoretical dual of the Voronoi diagram, the *Delaunay triangulation*, denoted $DT(S)$, is also a fundamental construction in geometry [38]. Given a set of points S in a metric space $(\mathbb{X}, d)$, we want to construct a triangulation of S such that we maximize the lexicographical order of its angles. As it turns out this graph is exactly the dual of the Voronoi diagram of those points. Here two sites in S are connected if and only if their Voronoi cells border each other (see Figure 2.2(b)).

Definition 2.4.4 (Delaunay Triangulation). Given a set of sites S , the Delaunay triangulation is the graph (S, E) , where for $x, y \in S$, (x, y) is an edge in E if and only if $V(x) \cap V(y) \neq \emptyset$.

Delaunay triangulations have a large array of applications, and we again refer the reader to Aurenhammer's survey [36].

2.5 Conical Funk

While geometries like the Hilbert and Funk are commonly defined in terms of bounded convex bodies, there are applications in analysis where it is of particular interest to examine unbounded convex sets, particularly cones [22]. When we refer to a cone $C \subset \mathbb{R}^d$, we mean a fully dimensional convex set which is closed under positive scaling (that is, whenever $p \in C$ so is βp for any $\beta \geq 0$). We assume our cones are *pointed*, meaning that they contain no lines. This implies that each cone has a unique *apex* at the origin. We denote the boundary of C by ∂C and its interior by $\text{int } C$. We will examine two types of cones, elliptical and polygonal cones (see Figure 2.3). For a point $p \in \mathbb{R}^d$, we write C_p for the cone C translated so that its apex is at p and C_p^r for the cone $-C$ (obtained by reflecting C) translated so that its apex is at p .




Figure 2.3: (a) An elliptical cone. (b) A polygonal cone.

Cones come with a natural partial order which is used to define the Funk metrics.

Definition 2.5.1 (Conical Partial Order). Given a convex cone C , we define the following partial order on $\text{int } C$: for $a, b \in \text{int } C$, we say $a \leq_C b$ if and only if $b - a \in C$. Equivalently, $a \leq_C b$ if and only if $b \in C_a$ (see Figure 2.4).




Figure 2.4: (a) A collection of 5 points in a cone C . (b) The cone C translated to the points. (c) A directed Hasse diagram describing their partial order, $d \leq_C e, c \leq_C e, b \leq_C c, a \leq_C b$.

To illustrate this partial order, we provide a Hasse diagram of a set of sites in the cone in Figure 2.4(c). We can define the Funk and reverse Funk conical metrics in a cone through the use of this partial order. The definition of the Funk metric here corresponds to the definition of the Funk metric in bounded convex bodies on bounded cross-sections of the cone [22].

Definition 2.5.2 (Forward Funk metric (conical)). For $a, b \in \text{int } C$, the forward Funk metric $F_C(a, b)$ is

$$F_C(a, b) = \ln \inf\{\beta \in \mathbb{R} \mid a \leq_C \beta b\}.$$

Definition 2.5.3 (Reverse Funk metric (conical)). For $a, b \in \text{int } C$, the reverse Funk metric is

$$F_C^r(a, b) = F_C(b, a).$$

Note that these definitions generalize the standard definition of the Funk metric in a bounded convex domain to a cone in a way that is consistent with the Hilbert metric definition on the cone [22]. In particular, for any $a, b \in \Omega$ where Ω is a bounded cross section of C , the restriction satisfies $F_\Omega(a, b) = F_C(a, b)$. The Funk metrics are weak metrics on C (see Definition 2.1.1) that satisfy neither the positivity property (property 1) nor the symmetry property (property 2). They only satisfy the triangle inequality (property 3).

In bounded convex domains, the Funk metric is always positive. In unbounded regions, however, numerators and denominators in the standard Funk formulation (see Remark 2.3 [22]) can be infinite. The cone-based definitions above (see Definitions 2.5.2 and 2.5.3), which follow Birkhoff's version of the Hilbert metric [22] handle this naturally. A consequence of these definitions is that the Funk metrics become signed distance functions and can take negative values. When b is in C_a , the value of $\ln(\beta)$ is negative; when b is on the boundary of C_a , $\ln(\beta)$ is 0; and when b is below C_a , $\ln(\beta)$ is positive (see Figure 2.5 (a)-(c)).




Figure 2.5: Two points a and b where (a) $F_C(a, b) < 0$, (b) $F_C(a, b) = 0$, (c) $F_C(a, b) > 0$.

Since the infimum definition of the forward Funk metric is not directly amenable to computation, we provide an alternative. To compute $F_C(a, b)$, it suffices to find the β such that βb intersects

∂C_a (see Figure 2.5). This geometric observation leads to an equivalent and more computationally feasible definition using similar triangles.

Proposition 2.5.4. *Let $a, b \in \text{int } C$. Consider the 2-dimensional slice through a, b , and the origin. Suppose the ray from the origin through b intersects ∂C_a on a side parallel to a ray R of C . Then, $F_C(a, b) = \ln \frac{d(a, R)}{d(b, R)}$, where $d(a, R)$ and $d(b, R)$ denote the perpendicular Euclidean distances from a and b to R respectively.*

Proof. In the 2-dimensional slice let βb be where the ray through the origin and b intersects ∂C_a . Since this boundary ray is parallel to R we can see that $d(\beta b, R) = d(a, R)$. By similar triangles we can see that $\frac{d(\beta b, R)}{d(b, R)} = \beta$. Hence $\frac{d(a, R)}{d(b, R)} = \beta$. Taking the natural log of both sides gives the result. $\square$

Recall that the ball of radius r around p is $B(p, r) = \{q : d(p, q) \leq r\}$. We now characterize balls in the forward and reverse Funk metrics.

Lemma 2.5.5. *Given $a \in \text{int } C$ and $r \in \mathbb{R}$, the forward and reverse Funk balls satisfy:*

$$(i) \quad B^F(a, r) = C_{e^{-r}a}$$

$$(ii) \quad B^{rF}(a, r) = C_{e^r a}.$$

Proof. We show that $p \in B^F(a, r)$ if and only if $p \in C_{e^{-r}a}$. Consider the 2D slice of the cone given by the plane through a, p , and the origin. By the similar triangles formula, $F_C(a, p) = \ln \frac{d(a, R)}{d(p, R)}$, where R is the appropriate ray of C . Since scaling a point by e^{-r} scales its distance to R by the same factor, we have $d(e^{-r}a, R) = e^{-r}d(a, R)$. Therefore,

$$F_C(a, p) \leq r \iff \frac{d(a, R)}{d(p, R)} \leq e^r \iff \frac{e^{-r}d(a, R)}{d(p, R)} \leq 1 \iff d(e^{-r}a, R) \leq d(p, R).$$

This holds if and only if $p \in C_{e^{-r_a}}$. The reverse Funk is analogous. □

Lemma 2.5.6. *Straight lines are geodesic in the forward and reverse Funk metrics.*

Proof. We first consider the forward Funk metric. Consider three collinear points $a, b, c \in \text{int } C$ where b lies on the line segment a to c . The ray from a through b is the same as the ray from a through c , so both distances $F_C(a, b)$ and $F_C(a, c)$ are computed using the same boundary ray R of C . Writing out the distances, we have

$$\begin{aligned} F_C(a, b) + F_C(b, c) &= \ln \frac{d(a, R)}{d(b, R)} + \ln \frac{d(b, R)}{d(c, R)} \\ &= \ln \frac{d(a, R)}{d(c, R)} = F_C(a, c). \end{aligned}$$

The reverse Funk is analogous. □

Chapter 3: Voronoi Diagrams in the Hilbert Metric

3.1 Introduction

As mentioned in the preliminaries, Voronoi diagrams are fundamental structures in computational geometry which encode proximity information. They can be defined in any metric. In this chapter, we study the geometric and combinatorial properties of the Voronoi diagram of a set of point sites under the Hilbert metric. Given any m -sided convex polygon Ω in the plane, we present two randomized incremental algorithms and one deterministic divide and conquer algorithm. The first randomized algorithm and the deterministic algorithm compute the Voronoi diagram of a set of n point sites. The second randomized algorithm extends this to compute the Voronoi diagram of the set of n sites, each of which may be a point or a line segment. Our randomized algorithms all run in expected time $O(mn \log n)$. The divide and conquer algorithm runs in time $O(mn \log n)$. The algorithms use $O(mn)$ storage, which matches the worst-case combinatorial complexity of the Voronoi diagram in the Hilbert metric.

3.2 Hilbert Metric Balls in Higher Dimensions

In this section we consider the properties of metric balls in the Hilbert metric. Given a convex body Ω in $\mathbb{R}^d$, $p \in \text{int}(\Omega)$, and $r \geq 0$, we denote the *Hilbert ball* of radius r centered at p by

$$B_\Omega(p, r) = \{q \in \Omega \mid H_\Omega(q, p) \leq r\}.$$

We can extend the notion of distance to any nonempty set P in the natural way by taking the minimum distance to the set, that is, $H_\Omega(q, P) = \inf\{H_\Omega(q, p) \mid p \in P\}$. This allows us to talk about the Hilbert ball of other shapes, such as line segments. We will consider balls generated by both points and line segments.

3.2.1 Hilbert Balls of Points

Nielsen and Shao [24] provided a characterization of Hilbert balls when Ω is an m -sided convex polygon in $\mathbb{R}^2$, showing that the ball is a convex polygon bounded by at most $2m$ sides. We generalize their result to arbitrary convex polytopes in $\mathbb{R}^d$. This provides an alternative (and more elementary) proof that Hilbert balls are convex (also proved in [35, 39]).

Lemma 3.2.1. *Given any convex polygon Ω in $\mathbb{R}^2$ bounded by m sides, $p \in \text{int}(\Omega)$, and $r \geq 0$, the ball $B_\Omega(p, r)$ is a convex polygon bounded by at most $2m$ sides. It can be constructed in $O(m)$ time.*

Proof. For now, let us take Ω to be any bounded convex body (not necessarily a polytope) in $\mathbb{R}^2$. Let $\mathbb{S}^1$ denote the unit sphere, that is the set of all unit vectors in $\mathbb{R}^2$. For each $u \in \mathbb{S}^1$, consider the ray emanating from p in direction u , and define $q_r(u) \in \Omega$ to be the (unique) point along this ray

whose Hilbert distance from p is r . Clearly, the boundary of $B_\Omega(p, r)$ is just the set of points $q_r(u)$ over all $u \in \mathbb{S}^1$.

For any $u \in \mathbb{S}^1$, let $h(u)$ denote any support line of Ω at the point where the ray emanating from p in the direction u intersects the boundary of Ω (see Figure 3.1(a)). Define $h(-u)$ analogously for the ray emanating from p in the direction $-u$. Let $g(u)$ ($= g(-u)$) denote the point where $h(u)$ and $h(-u)$ intersect. (Through an infinitesimal perturbation of Ω , we may assume that $h(u)$ and $h(-u)$ are not parallel.) Let $b_r(u)$ denote the ray originating at $g(u)$ and passing through $q_r(u)$, and define $b_r(-u)$ analogously for $q_r(-u)$.




Figure 3.1: The Hilbert ball of a point p .

Let $W(u)$ denote the convex wedge whose apex is at $g(u)$ and is bounded by $h(u)$ and $h(-u)$ (see Figure 3.1(b)). It follows from basic properties of projective geometry that every line passing through p on this plane is cut by $h(u)$, $h(-u)$, $b_r(u)$, and p into four points that share the same cross ratio. This implies that every point on $b_r(u)$ is at Hilbert distance r from p with respect to $W(u)$. The same applies symmetrically to $b_r(-u)$, and therefore the sub-wedge of $W(u)$ bounded by $b_r(u)$ and $b_r(-u)$ is just the Hilbert ball of radius r centered at p for $W(u)$. (This was observed by Nielsen and Shao in their analysis of the two-dimensional case.) Let us call this sub-wedge $B_r(u)$ (shaded in yellow in Fig. 3.1(b)).

Note that $p \in \Omega \subseteq B_r(u)$, and therefore Hilbert distances from p in $B_r(u)$ are at most as large as they are in Ω . Therefore, for all $u \in \mathbb{S}^1$, $B_\Omega(p, r) \subseteq B_r(u)$, and hence

$$B_\Omega(p, r) \subseteq \bigcap_{u \in \mathbb{S}^1} B_r(u).$$

On the other hand, by definition of $B_r(u)$, we know that $B_r(u)$ and $B_\Omega(p, r)$ both cover the exactly the same portion of the chord parallel to u passing through p . Since these chords cover all the boundary points of $B_\Omega(p, r)$, we have

$$B_\Omega(p, r) = \bigcap_{u \in \mathbb{S}^1} B_r(u).$$

Clearly, $B_r(u)$ is the intersection of a (possibly infinite) set of convex wedges, and so it is also convex.

Suppose now that Ω is an m -sided convex polygon. We say that two sides f^+ and f^- form a *complementary pair* if there exists $u \in \mathbb{S}^1$ such that the rays emanating from p in the directions u and $-u$ hit f^+ and f^- , respectively (see Figure 3.1(c)). We can partition the elements of $\mathbb{S}^1$ into equivalence classes according to the associated complementary pair. All the unit vectors u from any one equivalence class share the same support lines $h(u)$ and $h(-u)$, and therefore all of them contribute the same wedge $B_r(u)$ to $B_\Omega(p, r)$. By considering the chords $\chi(v, p)$ for each each of the m vertices of Ω , it follows directly that the number of complementary pairs is at most m . These chords partition Ω into at most m double wedges about p , each of which contributes two edges to the final ball, for a total of $2m$ sides.

The final Hilbert ball $B_\Omega(p, r)$ can be constructed in $O(m)$ time by observing that the sorted

sequence of double wedges about p can be generated in linear time and then generating the two line segments bounding $B_r(u)$ within each double wedge (assuming a standard representation of Ω as a cyclic sequence of vertices). $\square$

This result can be readily generalized to the higher dimensional case, but the worst-case number of bounding facets is quadratic in the number of facets.

Lemma 3.2.2. *Given any convex polytope Ω in $\mathbb{R}^d$ bounded by m facets, $p \in \text{int}(\Omega)$, and $r \geq 0$, $B_\Omega(p, r)$ is a convex polytope bounded by $O(m^2)$ facets.*

Proof. The proof of Lemma 3.2.1 can be generalized by applying the construction there in dimension d . Let $\mathbb{S}^{d-1}$ denote the unit sphere in $\mathbb{R}^d$, that is, the set of all unit vectors in d -dimensional space. Each pair u and $-u$ in $\mathbb{S}^{d-1}$ generates a pair of $(d-1)$ -dimensional supporting hyperplanes $h(u)$ and $h(-u)$ that intersect on a $(d-2)$ -dimensional affine subspace $g(u)$. This defines a convex wedge $W(u)$ bounded by these two hyperplanes. Again, for any $r \geq 0$, there is a sub-wedge $B_r(u)$ that defines the Hilbert ball of radius r about p with respect to the wedge, since all chords passing through p intersect the sides of these wedges at points sharing the same cross ratios. The final Hilbert ball results by intersecting all of these sub-wedges, resulting in a convex polytope.

We say that two facets f^+ and f^- form a *complementary pair* if there exists $u \in \mathbb{S}^{d-1}$ such that the rays emanating from p in the directions u and $-u$ hit f^+ and f^- , respectively. As in the proof of Lemma 3.2.1, we partition the elements of $\mathbb{S}^{d-1}$ into equivalence classes according to the associated complementary pair. Clearly, there are at most $\binom{m}{2} = O(m^2)$ complementary pairs, and hence $B_\Omega(p, r)$ is bounded by at most this many pairs of facets. (It is not difficult to create polytopes achieving this asymptotic bound.) $\square$

Given the characterization on Hilbert ball in $\Omega \in \mathbb{R}^2$, a natural question arises, how can we

define Hilbert balls for points at infinity, that is, on $\partial\Omega$? Let q be a point in Ω , we define the Hilbert ball at $p \in \partial\Omega$ whose boundary passes through q to be limit of any sequence of Hilbert balls passing through q whose centers approach p . We show this limit is unique and yields a convex polygon with at most m sides.

Lemma 3.2.3. *Given any convex polygon Ω in $\mathbb{R}^2$ bounded by m sides, $p \in \partial\Omega$, $q \in \text{int}(\Omega)$, there exists a unique ball, B , centered at p containing q on its boundary which is a convex polygon with at most m sides.*

Proof. Let us begin by considering an arbitrary sequence of Hilbert balls $\{B_n\}_{n \in \mathbb{N}}$ centered at points p_n such that $\{p_n\}_{n \in \mathbb{N}}$ converges to a point p on a side e of Ω (see Fig. 3.2(a)). Since our sequences converges to p there must be some $n_0 \in \mathbb{N}$ such that $\forall n \geq n_0$ in the sequence the wedges that contributes the boundary sides to $B_n, n \geq n_0$ all share the side e except the one made by the spokes through the vertices of e and p . Call this other wedge W with sides f^+ and f^- . It can be see that as $\{p_n\}_{n \geq n_0}$ approaches p , the measure of the region contributed by the wedge W vanishes to 0. Likewise, for any wedge W' between some e' and e , the portion of the boundary contributed between p_n and e vanishes in the limit. Hence we need only look at the boundaries contributed by the wedges between e and other sides of Ω .




Figure 3.2: The Hilbert ball of a point p on $\partial\Omega$ through a point q .

To construct this boundary let q lie on the wedge between e and e'' from the perspective of p_n . By tracing the line through q to the apex of the wedge we are able to determine the boundary of

an arbitrary p_n through q and a point on the boundaries of the neighboring wedges, from here we can walk along the wedges using the same process to complete the boundary. As p_n approaches p we can see that this converges to the unique polygonal boundary made by the same procedure applied to p . Since each wedge contributes one side to the boundary and there are $m - 1$, paired with the edge e this gives the ball centered at p has at most m sides. $\square$

3.2.2 Hilbert Balls of Lines and Line Segments

In this section we will describe the ball of radius r about an arbitrary line segment in Ω . Let us first consider the case of line ℓ that intersects Ω . Let $B_\Omega(\ell, r)$ denote the set of points of Ω that lie within Hilbert distance r of ℓ . We will consider just the 2-dimensional case, but the general case in $\mathbb{R}^d$ can be generated by considering the union of 2-dimensional slices generated by planes passing through ℓ . Given any point q in the interior of Ω , the following lemma characterizes the point of ℓ that is closest to q .

Lemma 3.2.4. *Given any convex body Ω in $\mathbb{R}^2$ in general position and a line ℓ that intersects the interior of Ω , for any $q \in \text{int}(\Omega)$ that is not on ℓ , its closest point on ℓ is uniquely determined to be the point $p \in \ell$ such that there exist support lines from each of the endpoints of the chord $\chi(q, p)$ that intersect at a point lying on ℓ (see Fig. 3.3(a)).*

Proof. We begin by showing that, assuming general position, for any point $q \in \text{int}(\Omega)$, there exists a unique point p satisfying the desired properties. We first consider the simpler case where Ω is strictly convex, and we will generalize later. Let us also assume for the sake of illustration that ℓ is horizontal. Define t_0 and t_1 to be the leftmost and rightmost points of intersection of ℓ with Ω , respectively (see Fig. 3.3(b)).




Figure 3.3: The closest point to q on line ℓ .

Consider a point t that moves continuously along $\ell \setminus \Omega$, which starts at t_0 , moves to the left until reaching $x = -\infty$, then wraps around to $x = +\infty$, and finally moves to t_1 . At each such point t , let $h^+(t)$ and $h^-(t)$ denote the two support lines of Ω that pass through t above and below ℓ , respectively. Since Ω is strictly convex, each pair of supporting lines defines a unique chord by joining the points where these support lines intersect $\partial \Omega$. It is a straightforward consequence of convexity that points where these support lines intersect $\partial \Omega$ move continuously and monotonically along both sides of the boundary of Ω from t_0 to t_1 . Thus, every point of Ω lies on exactly one of these chords. Given point $q \in \text{int}(\Omega)$, the chord passing through q satisfies the conditions of the lemma. (Note that if Ω has vertices, multiple points t may generate the same chord. So even though the chord is unique, the associated supporting lines need not be.)

We can generalize to the case where Ω is not strictly convex. Suppose, for example, that it is an m -sided convex polygon. We may assume through an infinitesimal perturbation of ℓ that the linear extensions of any two edges of Ω do not intersect on ℓ . If we extend each edge of Ω until it intersects ℓ , we have m distinct points on $\ell \setminus \Omega$. It is no longer true that each point t generates a unique chord, but each point along the extended edge generates a chord in conjunction with the supporting vertex on the opposite side ℓ (see Fig. 3.3(c)). It follows from monotonicity that for any $q \in \text{int}(\Omega)$, exactly one of these chords passes through q .

Because the endpoints of each of these chords lies on opposite sides of ℓ , each chord intersects ℓ . We assert that the point p where this chord intersects ℓ is the unique closest point to q . To see why, consider the wedge $W(t)$ bounded by $h^+(t)$ and $h^-(t)$. Because the three lines ℓ , $h^+(t)$, and $h^-(t)$ are coincident at t , every line passing through q cuts these lines so that the cutting points have the same cross ratios. It follows that every point on ℓ is equidistant to q with respect to the Hilbert distance defined by $W(t)$. But (assuming general position) Ω intersects each such line through q along a strictly smaller subsegment, and hence the Hilbert distance between q and any point on ℓ other than p is strictly larger than $H_\Omega(q, p)$. Therefore, p is the unique closest point to q on ℓ . $\square$

Lemma 3.2.5. *Given any convex polygon Ω in $\mathbb{R}^2$ bounded by m sides, any line ℓ that intersects the interior of Ω , and $r \geq 0$, the Hilbert ball $B_\Omega(\ell, r)$ is a convex polygon bounded by at most $2m$ sides. Furthermore, this ball can be constructed in $O(m)$ time.*

Proof. Our proof is constructive. As in Lemma 3.2.4, let us assume that ℓ is horizontal. Consider the m points where the linear extensions of the edges of Ω intersect ℓ . We assume by general position that these points are distinct. As mentioned in the earlier proof, each of these points t generates two support lines $h^+(t)$ and $h^-(t)$ lying above and below ℓ , respectively. It also generates a continuous family of chords, by joining each point on the corresponding side $e(t)$ of Ω to the vertex of Ω through which the other support line passes (see Fig. 3.4(a) and (b)).

Let $W(t)$ denote the double wedge bounded by $h^+(t)$ and $h^-(t)$. Given the radius r , we can draw a ray $b_r^+(t)$ emanating from t and passing above ℓ so that for every point q on $b_r^+(t)$, its Hilbert distance from ℓ with respect to $W(t)$ is r . As shown in Lemma 3.2.4, if q also lies on one of the chords generated by edge $e(t)$, then this is also true with respect to the Hilbert distance




Figure 3.4: The Hilbert ball of a line ℓ .

defined by Ω . Therefore, such a point q lies on the upper boundary of $B_\Omega(\ell, r)$ (see Fig. 3.4(b)). Analogously, we can find the ray $b_r^-(t)$, which generates the lower boundary of $B_\Omega(\ell, r)$ for the chords generated by $e(t)$.

All that remains to complete the entire ball is to repeat this process for each of the m points t where the extensions of the sides of Ω intersect ℓ . The ball $B_\Omega(\ell, r)$ is clearly a simple polygon bounded by at most $2m$ sides. Convexity follows by observing that, due to the monotonicity of the points t along the line ℓ , the slopes of its sides along the upper boundary of the ball decrease monotonically and the slopes of sides along the lower boundary of the ball increase monotonically. $\square$

Given two points $a, b \in \Omega$, the Hilbert ball of a line segment $\overline{ab}$ can be formed by first constructing the Hilbert ball for the line ℓ that passes through these points (see Fig. 3.5(a)), then cutting this ball through the chords of the line-ball construction passing through points a and b (see Fig. 3.5(b)), and finally filling in the missing ends with the portions of the Hilbert balls centered at the points a and b (see Fig. 3.5(c)).

Lemma 3.2.6. *Given any convex polygon Ω in $\mathbb{R}^2$ bounded by m sides, any line segment $\overline{ab}$ for $a, b \in \Omega$, and $r \geq 0$, the Hilbert ball $B_\Omega(\overline{ab}, r)$ is a convex polygon bounded by at most $4m$ sides. Furthermore, this ball can be constructed in $O(m)$ time.*




Figure 3.5: Three-step process for building the Hilbert ball of a line segment $\overline{ab}$.

3.3 Characterizing Voronoi Diagrams in the Hilbert Metric

Using our understanding of Hilbert balls we can characterize the Voronoi diagram of a set of point sites in the Hilbert metric. Throughout, let Ω denote a convex polygon in $\mathbb{R}^2$, and unless otherwise stated, distances will be in the Hilbert metric induced by Ω , which we denote by $H_\Omega(\cdot, \cdot)$ or simply $H(\cdot, \cdot)$, when Ω is clear. Let S denote a set of n points lying within the interior of Ω , which we call *sites*. For $p \in S$, define its *Voronoi cell* to be

$$V(p) = V_S(p) = \{q \in \Omega : H(q, p) \leq H(q, p'), \forall p' \in S \setminus \{p\}\}.$$

Although points on the boundary of Ω are infinitely far from points in the interior of Ω , we can compare the relative distances between a fixed boundary point and two interior points by considering the limit as an interior point approaches this boundary point. The *Voronoi diagram* of S in the Hilbert metric induced by Ω , denoted $\text{Vor}_\Omega(S)$, is the cell complex of Ω induced by the Voronoi cells $V(p)$ for all $p \in S$. We assume that the points of S are in general position, and in particular, the line passing through any pair of sites of S and the lines extending any two edges of Ω are not coincident at a common point (including all three being parallel). If this assumption does not hold, the bisectors separating Voronoi cells can widen into 2-dimensional regions.

Recall that a set $R \subseteq \mathbb{R}^d$ is a *star* (or is *star-shaped*) with respect to a point $p \in R$ if for each $q \in R$, the segment pq lies within R . We next show that Hilbert Voronoi cells are *stars*.

Lemma 3.3.1. *Voronoi cells in the Hilbert Metric are stars with respect to their defining sites.*

Proof. If this were not the case, there would exist a site p and points $x, y \in \Omega$ such that $x \in V(p)$, $y \notin V(p)$, and y lies on the line segment px . By collinearity, $H(p, x) = H(p, y) + H(y, x)$. Letting q be the closest site to y , we have $H(q, y) < H(p, y)$ and $H(p, x) \leq H(q, x)$. Combining these we have $H(q, x) + H(p, y) > H(q, y) + H(p, x)$, or equivalently $H(q, x) > H(q, y) + H(y, x)$. But this violates the triangle inequality, yielding a contradiction. $\square$

3.4 Bisectors in the Hilbert Metric

Given two sites $p, p' \in \text{int}(\Omega)$, we define their Hilbert bisector, denoted the (p, p') -bisector, to be $\{z \in \Omega : H_\Omega(z, p) = H_\Omega(z, p')\}$. We will explore the conditions for a point z to lie on the bisector. Let x and y denote the endpoints of the chord $\chi(z, p)$, and define x' and y' analogously for $\chi(z, p')$. Label these points in the order $\langle x, z, p, y \rangle$ and $\langle x', z, p', y' \rangle$ (see Figure 3.6(a)). Finally, let ℓ_x, ℓ_p , and ℓ_y denote the lines passing through the line segments xx', pp' and yy' , respectively. If these three lines are coincident on some point q then by basic properties of projective geometry, the cross ratios $(z, p; y, x)$ and $(z, p'; y', x')$ are equal. It follows that $H_\Omega(z, p) = H_\Omega(z, p')$, and hence z is on the bisector. If not, then the cross ratios are different and the Hilbert distances are different. Observe that as z approaches the boundary of Ω (on the same side of ℓ_p as ℓ_x), the points z, x , and x' converge on a common point on $\partial\Omega$, and ℓ_x approaches a support line for Ω at this point (see Figure 3.6(b)). Thus, we obtain the following characterization of bisector points.




Figure 3.6: Conditions for a point z to lie on the Hilbert bisector between sites p and p' .

Lemma 3.4.1. *Given a convex body Ω in $\mathbb{R}^2$, sites $p, p' \in \text{int}(\Omega)$ and any other point $z \in \text{int}(\Omega)$, z lies on the (p, p') -bisector if and only if lines ℓ_x , ℓ_p , and ℓ_y (defined above) are coincident. Further, a point $x \in \partial\Omega$ is on the Hilbert bisector if the coincidence holds when ℓ_x is any support line at x .*

When Ω is an m -sided convex polygon in $\mathbb{R}^2$, we can provide a more precise characterization of the bisectors. Given two sites $p, p' \in \text{int}(\Omega)$, we will show below that the (p, p') -bisector is a piecewise conic (see Lemma 3.4.5). Assuming this for now, we can characterize the breakpoints in this curve. Letting $\{v_1, \dots, v_m\}$ denote the vertices of Ω , the $2m$ chords $\chi(v_i, p)$ subdivide Ω into $2m$ triangular regions, which we call p 's *sectors* with respect to Ω (see Figure 3.7(a)).




Figure 3.7: (a) The sectors of p with respect to Ω and (b) the Hilbert bisector (in red) between p and p' .

3.4.1 Bisector Segments and Combinatorial Complexity

For each point z on the (p, p') -bisector, let e and f denote the edges where the endpoints of chord $\chi(z, p)$ intersect $\partial\Omega$, and let e' and f' denote the corresponding edges of $\chi(z, p')$. Observe that the pair (e, f) is uniquely determined by the sector of p containing z , and the pair (e', f') is similarly determined by the sector of p' containing z . Therefore, the points of the (p, p') -bisector can be grouped into a discrete set of equivalence classes based on their sector memberships with respect to p and p' (see Figure 3.7(b)). We refer to these as *bisector segments*. Observe that by the star-shaped nature of Voronoi cells, as we travel along the bisector, we encounter the (up to $2m$) sectors of p in cyclic order (say clockwise as in the figure), and we visit (up to $2m$) the sectors of p' in the opposite cyclic order (counterclockwise). We will see below that, each bisector segment can be described by a simple parametric function involving p, p' , and the four edges defining the equivalence class. Star-shapedness implies that the bisector is simply connected. Combining the total number of sectors involved for each site, we have:

Lemma 3.4.2. *Given an m -sided convex polygon Ω in $\mathbb{R}^2$ and sites $p, p' \in \text{int}(\Omega)$, the (p, p') -bisector is a simply connected piecewise curve consisting of at most $4m$ bisector segments.*

Given a set of n sites S in Ω , the Voronoi diagram consists of a collection of n Voronoi cells. The intersection of two cells $V(p)$ and $V(p')$ (if nonempty) is a portion of (p, p') -bisector, which we call a *Voronoi edge*. As shown in the above lemma, each such edge is composed of at most $4m = O(m)$ bisector segments. A cell's boundary may also contain a portion of the boundary of Ω . The intersection of two Voronoi edges (if nonempty) is a *Voronoi vertex*. Because the diagram is a planar graph, we have the following bounds by a straightforward application of Euler's formula.

Lemma 3.4.3. *Given an m -sided convex polygon Ω in $\mathbb{R}^2$ and a set of n sites S in Ω , $\text{Vor}_\Omega(S)$ has n Voronoi cells, at most $3n$ Voronoi edges, and at most $2n$ Voronoi vertices. Each Voronoi edge consists of at most $4m$ bisector segments. Therefore, the entire diagram has total combinatorial complexity $O(mn)$. The average number of Voronoi edges per cell is $O(1)$, and the average number of bisector segments per cell is $O(m)$.*

The following lemma shows that the bound on the combinatorial complexity is tight in the worst case.

Lemma 3.4.4. *For all sufficiently large m and n , there exists a convex polygon Ω' with m sides and a set S of n sites within Ω' such that $\text{Vor}_{\Omega'}(S)$ has combinatorial complexity $\Omega(mn)$. (Here we are using $\Omega(\cdot)$ in the asymptotic sense.)*

Proof. We start the construction with an axis parallel rectangle R that is slightly taller than wide, and the set S consists of n points positioned on a short horizontal line segment near the center of this rectangle (see Figure 3.8(a)). (This violates our general position assumptions, but the construction works if we perturb the sides of the rectangle.) Also, create a set of $m - 4$ points $V = \{v_1, \dots, v_{m-4}\}$ along the left edge of R . We can adjust the spacing of points of V and S and the side lengths of R so that there exists a diamond shape (shaded in orange in the figure) so that for any point q in this region, the chord $\chi(v_i, q)$ intersects the boundary of R along its vertical sides. Observe that within the diamond, each consecutive pair of sites contributes an edge to the Voronoi diagram, which is easily verified to be a vertical line segment within the diamond.

To form Ω' , we bend the left side of the rectangle out infinitesimally (see Figure 3.8(b)) so the points of V become vertices of a convex polygon. The bending does not alter the shape of the Voronoi edges significantly, except to break each edge up into multiple arcs.




Figure 3.8: Proof of Lemma 3.4.4.

We assert that each consecutive pair of points $p, p' \in S$ contributes at least $m - 4$ arcs to the Voronoi diagram. To see why, observe that there are $m - 4$ sectors about p and p' , and vertical line passing between p and p' intersects the boundaries between these sectors (see the red curve in Figure 3.8(c)). Due to the infinitesimal bending of the left side of R , each crossing produces a new segment on the bisector. Since there are $n - 1$ consecutive pairs of sites, the total complexity of $\text{Vor}_{\Omega}(S)$ is at least $(m - 4)(n - 1) = \Omega(mn)$, as desired. $\square$

3.4.2 Bisector Segments Structure

In this section we discuss the properties of bisectors in the Hilbert metric. We consider cases involving both points, lines, and combinations thereof. In particular, we prove that the bisectors are piece-wise conics.

First, we will characterize the local structure of a bisector between two point sites a and b with respect to Ω . Let us hypothesize that a point $q \in \text{int}(\Omega)$ lies on the bisector between these sites (see Fig. 3.9(a)). Draw chords χ_{aq} and χ_{bq} . Let us make the general-position assumption that these chords do not intersect vertices of Ω . Consider the two edges of Ω incident chord χ_{aq} , and let p_a be the point (possibly at infinity) where the linear extensions of these edges intersect (see Fig. 3.9(b)). Define p_b analogously for the chord χ_{bq} .




Figure 3.9: Segments of a point-point bisector

If there are only two distinct edges involved (implying that $p_a = p_b$), the problem can be reduced to a simple 1-dimensional case by projection through the point where the extensions of the two edges meet. Otherwise, there are two wedges with apexes at the points p_a and p_b such that Ω is contained within their intersection. The shape resulting from the intersection of the two wedges may be a convex quadrilateral (as shown in Fig. 3.9(b)), but if two edges coincide, we obtain a triangle (see Fig. 3.9(c)). Also, this shape may be unbounded, which we treat as wrapping around the projective plane.

Shoot two rays originating at p_a passing through a and q , and let a' and x denote the respective points these rays hit on the opposite side of the quadrilateral (see Fig. 3.9(b)). Define b' and y analogously for rays originating at p_b through b . Let us parameterize points on these two edges so that -1 and 1 are the edge endpoints, and the points along the edge appear in the orders $\langle -1, a', x, 1 \rangle$ and $\langle -1, b', y, 1 \rangle$, respectively. We can now think of points $\{a', x, b', y\}$ as reals, where $-1 < a' \leq x < 1$ and $-1 < b' \leq y < 1$.

Lemma 3.4.5. *Given a convex polygon Ω and two sites $a, b \in \text{int}(\Omega)$, the bisector between a and b in the Hilbert metric is a piecewise conic curve.*

Proof. The proof of Lemma 3.2.1 can be generalized by applying the construction there in dimen-

sion d . Let $\mathbb{S}^{d-1}$ denote the unit sphere in $\mathbb{R}^d$, that is, the set of all unit vectors in d -dimensional space. Each pair u and $-u$ in $\mathbb{S}^{d-1}$ generates a pair of $(d-1)$ -dimensional supporting hyperplanes $h(u)$ and $h(-u)$ that intersect on a $(d-2)$ -dimensional affine subspace $g(u)$. This defines a convex wedge $W(u)$ bounded by these two hyperplanes. Again, for any $r \geq 0$, there is a sub-wedge $B_r(u)$ that defines the Hilbert ball of radius r about p with respect to the wedge, since all chords passing through p intersect the sides of these wedges at points sharing the same cross ratios. The final Hilbert ball results by intersecting all of these sub-wedges, resulting in a convex polytope.

We say that two facets f^+ and f^- form a *complementary pair* if there exists $u \in \mathbb{S}^{d-1}$ such that the rays emanating from p in the directions u and $-u$ hit f^+ and f^- , respectively. As in the proof of Lemma 3.2.1, we partition the elements of $\mathbb{S}^{d-1}$ into equivalence classes according to the associated complementary pair. Clearly, there are at most $\binom{m}{2} = O(m^2)$ complementary pairs, and hence $B_\Omega(p, r)$ is bounded by at most this many pairs of facets. (It is not difficult to create polytopes achieving this asymptotic bound.) $\square$

Next, let us consider the case where a is a point site and $b = \overline{b_0 b_1}$ is a line segment, where $a, b_0, b_1 \in \Omega$. As before, we wish to characterize the set of points q lying on the Hilbert bisector between a and b with respect to Ω . Let ℓ_b denote the linear extension of this line segment.

Recall from Lemma 3.2.4 that the closest point on ℓ_b is uniquely determined to be the point b_q such that there exist support lines to Ω at the endpoints of the chord $\chi(q, b_q)$ that intersect at a point on ℓ_b (see Fig. 3.10(a)).

Given the point p_b , we carry out the same construction as in the point-point bisector case, where the point p_a is defined exactly as before. We centrally project the points a and q through p_a until they hit the far side of the p_b -wedge, and symmetrically we centrally project the segment

3.5 Randomized Incremental Algorithm

In this section we present two randomized incremental algorithms to compute the Hilbert Voronoi diagram of a set S of n sites. In both cases, the domain is an m -sided convex polygon Ω . First, we describe the algorithm in the case when S of n point sites. Second, we explain the changes when the sites are line segments. Both algorithms follow the structure of the randomized incremental algorithm for abstract Voronoi diagrams by Mehlhorn, Meiser, and Ó'Dúnlaing [40].

3.5.1 The Hilbert Metric and Abstract Voronoi Diagrams

There are two important considerations when applying the techniques of Mehlhorn, Meiser, and Ó'Dúnlaing [40] to generate Voronoi diagrams in the Hilbert metric. First, given its unique geometry, the Hilbert metric does not inherently satisfy the properties of abstract Voronoi diagrams, though it does so in actuality. The distinction is that the Voronoi diagrams in the Hilbert metric don't partition $\mathbb{R}^2$, but instead partition Ω , which serves as a substitute. Second, out of the four operations that Mehlhorn, Meiser, and Ó'Dúnlaing posit work in constant time, one, finding the intersection of two bisectors runs in time $O(m)$. This final difference contributes a factor of m to our overall run-time.

3.5.2 Point Sites

In this section we consider the construction of a Voronoi diagram for a set S of n point sites. The algorithm given in [40] starts by computing a bounding enclosure, but our body Ω serves that role. Let S be our set of sites. Consider an arbitrary iteration of the algorithm, and let R denote the set of sites already added to the diagram. To facilitate the constructing of our algorithm we

maintain two data structures: $\text{Vor}(R)$ the current Voronoi diagram of R as a cell complex, and $G(R)$ the *conflict graph* of R with respect to S (defined below).

$\text{Vor}(R)$ may be stored in any standard data structure for planar subdivisions, such as a doubly connected edge list [38]. The vertices of $\text{Vor}(R)$ are the Voronoi vertices and the edges are the bisectors connecting them. To facilitate efficient tracing of bisectors, we augment each Voronoi cell with *spokes*, that is, line segments emanating from each site in the diagram to the segment vertices of the bisectors forming its Voronoi cell (see Fig. 3.11).

The conflict graph $G(R)$ is defined as follows. Its vertices consist of the edges of $\text{Vor}(R)$ together with the remaining sites $S \setminus R$. There is an edge between $e \in \text{Vor}(R)$ and $p \in S \setminus R$ if and only if the addition of p would result in edge e either being removed or trimmed.




Figure 3.11: Inserting a new site in the randomized algorithm.

The algorithm randomly permutes the sites and inserts them one by one. The first two sites are inserted by brute force and the conflict graph is initialized based on the single bisector between these sites. Otherwise, let us assume that we have already inserted some subset R of sites and are now inserting the next site $p \in S \setminus R$. We need to (1) trace the boundary of p 's Voronoi cell in $\text{Vor}(R \cup \{p\})$ and (2) update the conflict graph. To trace a bisector between two sites we use the following method:

Let $V(p)$ denote the Voronoi cell of p after its insertion. We select any edge e that conflicts

with p and we compute any point t on this edge that lies within p 's Voronoi cell. We proceed to walk along the bisector, continuously maintaining the distances to the associated sites until reaching a point on the boundary of $V(p)$. Starting at this point, we proceed to trace the boundary of $V(p)$ in counterclockwise order about p (see Figure 3.11(b)). This involves two types of traces:

Bisector Trace: We are tracing the bisector between p and some existing site p' (see Figure 3.11(c)).

We consider the sectors of p and p' containing the current portion of the (p, p') -bisector. By applying the parameterization from Lemma 3.4.5, we can trace the bisector until (1) we encounter the boundary of either sector, (2) we encounter the bisector between p' and another site, or (3) we encounter the boundary of Ω .

In the first case, we create a new segment vertex here, add spokes to this vertex, erase the extension of the sector edge (shown as a broken line in Figure 3.11(c)) and continue the tracing in the new sector. In the second case, we create a new Voronoi vertex, add spokes to this vertex from its three defining sites, and continue the trace along the new bisector. In the third case, we insert two spokes joining the point where the bisector encounters the boundary to p and p' , respectively. We then transition to the boundary trace described next.

Boundary Trace: We are tracing the Voronoi cell of p along the boundary of Ω . We walk along the boundary of Ω in counterclockwise order, considering each consecutive sector of the closest site p' , prior to p 's insertion. By applying the parameterization from Lemma 3.4.5, we determine whether the (p, p') -bisector intersects the boundary of Ω within the intersection of the current pair of sectors. If so, we identify this point on the $\partial\Omega$, add spokes to each of p and p' , and then resume tracing along the (p, p') -bisector. Otherwise, we encounter one of the two sector boundaries, and we continue the tracing the next sector.

Along the way, we erase spokes from the current diagram, and/or introduce new spokes connected to p . When we return to our starting point, the insertion is completed.

While we are tracing the bisector for the new site, we can simultaneously update the conflict graph. Whenever we arrive at a new Voronoi vertex in the diagram, we construct a Hilbert ball at this point whose radius is the distance to the newly added site p (recall Lemma 3.2.1). This includes Voronoi vertices “at infinity”, which arise whenever a bisector extends all the way to the boundary of Ω . Let p' and p'' be the two other sites of R that share this vertex. All the sites of $S \setminus R$ in this ball are added as conflict-graph neighbors of the two new Voronoi edges between p and p' and p and p'' .

The algorithm given in [40] runs in $O(n \log n)$ expected time, under the *update assumption* of Clarkson and Shor [41], which states that the time to insert a new site p is proportional to the number of objects that conflict with regions that conflict with p . In our context, we need to add an additional factor of $O(m)$. First, this is needed because bisectors are no longer of constant complexity, but consist of $O(m)$ segments. Our algorithm processes each new segment in $O(1)$ time. Second, when updating the conflict graph, in order to determine which points conflict with the newly created edges, we need to construct a Hilbert ball at each newly created Voronoi vertex, and this takes time $O(m)$. The overall expected running time is thus larger by a factor of $O(m)$, implying the following.

Theorem 3.5.1. *Given an m -sided convex polygon Ω in $\mathbb{R}^2$ and a set of n point sites S in Ω , the randomized incremental algorithm computes $\text{Vor}_\Omega(S)$ in expected time $O(mn \log n)$ (where the expectation is over all possible insertion orders).*

3.5.3 Segment Sites

In this section, we present the modifications necessary to generalize the algorithm from Section 3.5.2 to the case of line segments. Since the algorithm by Mehlhorn, Meiser, and Ó'Dúnlaing [40] is quite general, the approach applies here as well, but with a few changes.

The first change involves the use of *spokes* in the representation of the diagram. In the point-site case, each spoke was a line segment connecting a vertex v of on a bisector (which may be either a Voronoi vertex or the endpoint between bisector segments) to each of its defining sites. When one of these sites is a line segment, we join the vertex to its closest point on the line segment. This closest point is determined by the local minimality conditions given in Section 3.2.2. These spokes cannot intersect each other except possibly at their endpoints. This is a direct consequence of the triangle inequality and the fact that straight line segments are geodesics.

The second change is that bisector tracing involves a different procedure for the line-line and point-line cases, as opposed to the point-point case. However, as observed in Lemma 3.4.6, the sectors are still conics, and their worst-case complexity is no different than in the point-point case. There is one additional situation that arises when tracing these bisectors. When the closest part of the line segment is the segment's endpoint, we treat the site as if it was a point site. As the bisector moves around this endpoint, eventually the closest point transitions to the interior of the segment. This transition can be determined in $O(1)$ time (again, by using the local minimality conditions given in Section 3.2.2). When this happens, we transition to treating the site as if it was a line.

The final change is how the conflict graph is updated following each insertion. Recall that in the case of point sites, this was handled by computing the Hilbert ball centered at each newly created Voronoi vertex, and updating the conflict neighborhoods for all the points lying within this

ball. The process is exactly analogous here, but there may be many such balls associated with any given segment. As we are tracing the bisector bounding the Voronoi cell of the new segment, we may generally encounter multiple instances where we create a new Voronoi vertex involving an interior point on this line segment and two other sites (which might be points or line segments). The change is that all of the yet uninserted sites in the union of these balls need to have the conflict neighbors updated.

Other than these changes, the algorithm and its analysis go through exactly as before. Again, we suffer an additional factor of $O(m)$ in the running time due to addition complexity of the diagram and the complexity to compute Hilbert balls.

Theorem 3.5.2. *Given an m -sided convex polygon Ω in $\mathbb{R}^2$ and a set of n point and line-segment sites S in Ω , the randomized incremental algorithm computes $\text{Vor}_\Omega(S)$ in expected time $O(mn \log n)$ (where the expectation is over all possible insertion orders).*

3.6 Divide and Conquer Algorithm

In this section we present a divide-and-conquer algorithm for constructing the Hilbert Voronoi diagram for a set of n point sites contained in a m -sided convex polygon Ω . Our algorithm follows the structure of the standard divide and conquer algorithm for the Euclidean Voronoi diagram (see, e.g., [42]).

The algorithm begins by partitioning the sites about a vertical line ℓ into two subsets, S_L and S_R , of roughly equal sizes that lie to the left and right of this line, respectively. We recursively compute the Hilbert Voronoi diagrams $\text{Vor}_\Omega(S_L)$ and $\text{Vor}_\Omega(S_R)$ (see Figure 3.12(b) and (c)). To merge them, we compute the bisector between S_L and S_R , which we denote by $B(S_L, S_R)$. This

consists of those points that are equidistant (in the Hilbert sense) from their closest sites in S_L and S_R . We shall see below that this consists of a collection of curves, called *components*, each being the concatenation of Voronoi edges between two sites, one in S_L and one in S_R (see Figure 3.12(d)). Each component terminates on the boundary Ω (that is, no component is a closed loop). We will show below that it is possible to construct all the components of $B(S_L, S_R)$ by a single bottom-up traversal, which runs in $O(mn)$ time. Once the bisector is constructed, we complete the merging process, by discarding the portions of $\text{Vor}_\Omega(S_L)$ that lie on the right side of the bisector and the portions of $\text{Vor}_\Omega(S_R)$ that lie on the left side of bisector. Before describing the bisector-construction process, we explain the bisector's structure.




Figure 3.12: Divide-and-conquer algorithm, bisectors, and the proof of Lemma 3.6.1(iii).

Lemma 3.6.1. *Given two sets of sites S_L and S_R separated by a vertical line ℓ , the Hilbert bisector $B(S_L, S_R)$ consists of a collection of pairwise disjoint curves, called components, where:*

- (i) *Each component is the concatenation of Voronoi edges from the final diagram, each between a site in S_L and another in S_R .*
- (ii) *Every point on $B(S_L, S_R)$ is visible to some point on $\ell \cap \Omega$, that is, for each $x \in B(S_L, S_R)$ there exists a point $y \in \ell \cap \Omega$ such that the open line segment xy does not intersect $B(S_L, S_R)$.*
- (iii) *Each component is terminated by endpoints that lie on the boundary of Ω (and hence no*

component is a closed loop).

(iv) *There are at most n components, and the total combinatorial complexity of $B(S_L, S_R)$ is $O(mn)$.*

Proof. Claim (i) follows from basic facts about Voronoi diagrams. Every point $q \in B(S_L, S_R)$ is on the Voronoi cell of at least two sites, one on each side of ℓ . By Lemma 3.3.1, the line segment from q to each site lies entirely within the Voronoi cell, and hence does not intersect $B(S_L, S_R)$. One of these two lines must cross $\ell \cap \Omega$, which establishes (ii).

To prove (iii), suppose to the contrary that a component does not contact the boundary of Ω , implying that it forms a closed loop. Let q be the point of this loop that is farthest from ℓ (that is, has the largest x -coordinate). By claim (ii) and consideration points infinitesimally close on either side of q on the loop, q must be visible to two sites $p, p' \in S_L$ such that the segment qp passes above the loop and qp' passes below the loop (see Figure 3.12(e)). Therefore, q is a Voronoi vertex between p, p' , and the closest site p'' in S_R . Thus, q is the center of a Hilbert ball passing through these three sites. However, because the line segment pp' lies entirely to the left of ℓ and p'' lies between these points to the right of ℓ , this violates the convexity of Hilbert balls, a contradiction.

Claim (iv) follows from the fact that each component contains at least one full Voronoi edge from the final diagram, and by Lemma 3.4.3 there are at most $O(n)$ Voronoi edges in the final diagram. Also, the bisector is part of the final diagram, and hence its total complexity cannot exceed that of the final diagram, which by Lemma 3.4.3 is $O(mn)$. $\square$

Let us consider how the bisector $B(S_L, S_R)$ is constructed. As in the previous section, we will triangulate each Voronoi cell by adding *spokes*, that is, line segments between each site and the various vertices (Voronoi vertices, bisector segment endpoints, vertices of Ω) that lie on the

boundary of its Voronoi cell (see Figure 3.13(a)). As mentioned above, we construct the bisector from bottom to top. (Note, however, that the components of $B(S_L, S_R)$ need not be vertically monotone.) We start at the lowest endpoint of q of $\ell \cap \Omega$. Let p and p' be the sites of S_L and S_R that are closest to q , respectively. We may assume that the construction algorithms for $\text{Vor}_\Omega(S_L)$ and $\text{Vor}_\Omega(S_R)$ provide us with the identities of p and p' as well as the triangles of their respective Voronoi cells that contain q . In $O(1)$ time, we can determine the closer of p or p' to q . If it is p , we begin a boundary trace counterclockwise about $\partial \Omega$ and if it is p' , we begin a boundary trace clockwise about $\partial \Omega$. These boundary traces are described below.




Figure 3.13: Tracing $B(S_L, S_R)$. (Drawing is not geometrically accurate.)

Generally there are two types tracings that we need to perform, depending on whether we are tracing a component of $B(S_L, S_R)$ or walking along the boundary of Ω in search of the next component. These are similar to the tracings we described in the case of the randomized algorithm. Note that generally, we are tracing counterclockwise about the sites of S_L and clockwise about the sites of S_R (see Fig. 3.13).

Bisector Trace: We are tracing the bisector between two sites $p_i \in S_L$ and $p_k \in S_R$ (see Figure 3.11(c)). We consider the sectors of p_i and p_k containing the current portion of the (p_i, p_k) -bisector. As in the bisector tracing for the randomized algorithm, we employ Lemma 3.4.5

to trace the bisector until either (1) it intersects one of the two sector edges, (2) hits another bisector, or (3) hits the boundary of Ω .

In the first case, we erase the extension of the sector edge (shown as a broken line in Figure 3.13(d)), create a new segment vertex here, add spokes to this vertex, and continue the tracing in the new sector. In the second case, we create a new Voronoi vertex, add spokes to this vertex from its three defining sites, and continue the trace along the appropriate bisector (always walking on a bisector between a left-side site and a right-side site). In the third case, we insert two spokes joining the point where the bisector encounters the boundary to p_i and p_k , respectively. We then transition to the boundary trace described next.

Boundary Trace: There are two cases, depending on whether we are tracing along the boundary of a cell of S_L or S_R . In the former case, we walk counterclockwise along the boundary of Ω until reaching the end of the current sector, and in the latter we walk clockwise. We consider each consecutive sector of the closest sites p_i from S_L and p_j from S_R . By applying the parameterization from Lemma 3.4.5, we determine whether the (p_i, p_k) -bisector intersects the boundary of Ω within the intersection of the current pair of sectors. If so, we identify this point on the $\partial\Omega$, add spokes to each of p_i and p_k , and then resume tracing along the (p_i, p_k) -bisector (see Figure 3.13(b)–(d)). Otherwise, we encounter one of the two sector boundaries, and we continue the tracing along the boundary.

The tracing process runs in $O(1)$ time for each newly added segment to the bisector, for a total of $O(mn)$ time. The overhead for the divide-and-conquer is dominated by the tracing time, which implies that divide-and-conquer algorithm's overall running time satisfies the recurrence $T(n) = 2T(n/2) + mn$. This solves to $O(mn \log n)$. In summary, we obtain the following result.

Theorem 3.6.2. *Given an m -sided convex polygon Ω in $\mathbb{R}^2$ and a set of n point sites S in Ω , the divide-and-conquer algorithm computes $\text{Vor}_\Omega(S)$ in time $O(mn \log n)$.*

Chapter 4: Analysis of Dynamic Voronoi Diagrams in the Hilbert Metric

4.1 Introduction

In this chapter, we expand upon the results of the previous chapter. We study the properties of Hilbert bisectors and analyze Hilbert Voronoi diagrams in the dynamic setting. We provide explicit equations for Hilbert bisectors as conics, the types of conics they take on, and when they degenerate into 2-dimensional regions.

4.2 Hilbert Bisector analysis

Before we proceed we recall the results of the previous chapter that Voronoi cells in the Hilbert Metric are stars and with complexity $\Omega(mn)$.

It will be helpful to recall two terms, *spokes* and *sectors*. Given our convex body Ω and a site s in Ω it can be seen that the Hilbert distance from our site s to a point p is dependent on the edges of Ω that intersect $\chi(s, p)$ (recall that $\chi(s, p)$ denotes the chord through s and p in Ω). This is characterized by taking the vertices of Ω , which we denote by V , and partitioning Ω into cells determined by the chords $\chi(v, s)$ for all $v \in V$. We define a *spoke* $r_\Omega(v, s)$, $r_\Omega(s, v)$ to be the intersection of ray vs and sv respectively with Ω when v is a vertex of V . Nielsen and Shao showed that Hilbert balls are $\Theta(m)$ -sided polygons whose boundaries are constructed around these




Figure 4.1: Illustration of different *sectors*.

spokes [24].

These *spokes* subdivide Ω into polygonal regions, which we will call *sectors*. Each sector T has the property that for all points $p \in T$, given two sites s and t there are four edges (not necessarily unique) of $\partial\Omega$, say E_A, E_B, E_C, E_D , such that $\chi(s, p)$ always intersects Ω at E_A and E_B and $\chi(t, p)$ at E_C and E_D . We will write $S(s, t, E_A, E_B, E_C, E_D)$ to be the unique *sector* such that for any $p \in S(s, t, E_A, E_B, E_C, E_D)$, we have $\chi(s, p)$ has in order of intersection E_A, s, p, E_B and $\chi(t, p)$ has E_C, t, p, E_D (see Figure 4.1). As a consequence of this, we can see that the bisector between two sites, s and t , is composed of curves which depend piece-wise on the sectors they pass through.

Theorem 4.2.1. *The equation of the bisector between two sites in any particular sector is of the form $Ax^2 + Bxy + Cy^2 + Dx + Ey + F = 0$, with coefficients depending only on the line equations of the four edges and the sites.*

Proof. Suppose we are in some arbitrary sector $S(s, t, E_A, E_B, E_C, E_D)$. Let the equation of the edges be $a_1x + a_2y + a_3 = 0$, $b_1x + b_2y + b_3 = 0$, $c_1x + c_2y + c_3 = 0$, and $d_1x + d_2y + d_3 = 0$,

for E_A , E_B , E_C , and E_D respectively.

Recall the Euclidean distance between a point (x, y) and a line $ux + vy + l = 0$ is given by $D(u, v, l, x, y) = \frac{|ux+vy+l|}{\sqrt{u^2+v^2}}$. Let $p \in T$ be an arbitrary point with coordinates (p_x, p_y) . Notice that, without loss of generality, if q is the point where chord $\chi(p, s)$ hits $\partial\Omega$ towards s , then the ratio $\frac{\|q-p\|}{\|q-s\|}$ can be replaced with the ratio $\frac{D(a_1, a_2, a_3, p_x, p_y)}{D(a_1, a_2, a_3, s_x, s_y)}$ by similar triangles. Using this technique we set $H(s, p) = H(t, p)$ and solve, yielding:

$$\frac{D(a_1, a_2, a_3, p_x, p_y)}{D(a_1, a_2, a_3, s_x, s_y)} \frac{D(b_1, b_2, b_3, s_x, s_y)}{D(b_1, b_2, b_3, p_x, p_y)} = \frac{D(c_1, c_2, c_3, p_x, p_y)}{D(c_1, c_2, c_3, t_x, t_y)} \frac{D(d_1, d_2, d_3, t_x, t_y)}{D(d_1, d_2, d_3, p_x, p_y)}$$

The equation of the line $ux + vy + l = 0$ divides the plane in two, characterized by the sign of the function $g(x, y) = ux + vy + l$. Given this, since all our points are on the same side of each line, when we substitute in for the function D , we can remove the absolute value bars from our equations. It follows that:

$$\frac{a_1 p_x + a_2 p_y + a_3}{a_1 s_x + a_2 s_y + a_3} \frac{b_1 s_x + b_2 s_y + b_3}{b_1 t_x + b_2 t_y + b_3} = \frac{c_1 p_x + c_2 p_y + c_3}{c_1 t_x + c_2 t_y + c_3} \frac{d_1 t_x + d_2 t_y + d_3}{d_1 p_x + d_2 p_y + d_3}$$

Collecting constants on one side gives us:

$$\frac{b_1 s_x + b_2 s_y + b_3}{d_1 t_x + d_2 t_y + d_3} \frac{c_1 t_x + c_2 t_y + c_3}{a_1 s_x + a_2 s_y + a_3} = \frac{c_1 p_x + c_2 p_y + c_3}{a_1 p_x + a_2 p_y + a_3} \frac{b_1 p_x + b_2 p_y + b_3}{d_1 p_x + d_2 p_y + d_3}$$

Letting the left side be a constant k , we find that the bisector satisfies the algebraic equation:

$$Ap_x^2 + Bp_x p_y + Cp_y^2 + Dp_x + Ep_y + F = 0,$$

where the coefficients are:

$$A = b_1c_1 - a_1d_1k$$

$$B = b_2c_1 + b_1c_2 - a_1d_2k - a_2d_1k$$

$$C = b_2c_2 - a_2d_2k$$

$$D = b_3c_1 + c_3b_1 - a_3d_1k - a_1d_3k$$

$$E = b_3c_2 + b_2c_3 - a_2d_3k - a_3d_2k$$

$$F = b_3c_3 - a_3d_3k.$$

By the symmetry of these equations we can see that the bisector in $S(s, t, E_A, E_B, E_C, E_D)$ is the same as the bisector in $S(s, t, E_B, E_A, E_D, E_C)$.¹ □

We present several ways to simplify the bisector depending on the sector case.

4.2.1 Finding Bisectors in a Sector

There are three main sector types: sectors with four distinct edges, sectors with three, and sectors with two. We begin with the four edge case. It is a well known fact from projective geometry that there exists a projective transformation that maps any convex quadrilateral to another (see Figure 4.2). For the sake of completeness, we present the derivation of this transformation for the special case where the target is the unit square.

Observation 4.2.2. *Bisectors in the four-edge case can be projectively transformed and solved in the unit square.*

¹We would like to thank Daniel Skora for comments on a previous version of this proof.




Figure 4.2: A projection between two quadrilaterals.

Proof. Suppose we are in a strictly four-edge sector $S(s, t, E_A, E_B, E_C, E_D)$ with edges E_A, E_C, E_B, E_D counter clockwise. Let our edges extend out infinity and let $p_1 = E_A \cap E_D$, $p_2 = E_A \cap E_C$, $p_3 = E_B \cap E_C$, and $p_4 = E_B \cap E_D$. Consider the vector:

$$Q^T = [q_{1,x}, q_{1,y}, q_{2,x}, q_{2,y}, q_{3,x}, q_{3,y}, q_{4,x}, q_{4,y}],$$

where $q_1 = (0, 0)$, $q_2 = (0, 1)$, $q_3 = (1, 1)$, and $q_4 = (1, 0)$. Then we can calculate the projective transformation matrix T such that $P' = TP$ when T sends points P on our quadrilateral to the unit square. We force the right corner of T to be 1 to fix our scaling factor.

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} \begin{bmatrix} t_{11} & t_{12} & t_{13} \\ t_{21} & t_{22} & t_{23} \\ t_{31} & t_{32} & 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

We solve for the elements of T using the following matrices: $MV = Q$ where $V =$

$[t_{11}, t_{12}, \dots]^\top$, $Q = [q_{1,x}, q_{1,y}, \dots]^\top$ and M is:

$$\begin{bmatrix} p_{1,x} & p_{1,y} & 1 & 0 & 0 & 0 & -p_{1,x} q_{1,x} & -p_{1,y} q_{1,x} \\ 0 & 0 & 0 & p_{1,x} & p_{1,y} & 1 & -p_{1,x} q_{1,y} & -p_{1,y} q_{1,y} \\ p_{2,x} & p_{2,y} & 1 & 0 & 0 & 0 & -p_{2,x} q_{2,x} & -p_{2,y} q_{2,x} \\ 0 & 0 & 0 & p_{2,x} & p_{2,y} & 1 & -p_{2,x} q_{2,y} & -p_{2,y} q_{2,y} \\ p_{3,x} & p_{3,y} & 1 & 0 & 0 & 0 & -p_{3,x} q_{3,x} & -p_{3,y} q_{3,x} \\ 0 & 0 & 0 & p_{3,x} & p_{3,y} & 1 & -p_{3,x} q_{3,y} & -p_{3,y} q_{3,y} \\ p_{4,x} & p_{4,y} & 1 & 0 & 0 & 0 & -p_{4,x} q_{4,x} & -p_{4,y} q_{4,x} \\ 0 & 0 & 0 & p_{4,x} & p_{4,y} & 1 & -p_{4,x} q_{4,y} & -p_{4,y} q_{4,y} \end{bmatrix}$$

Since the Hilbert metric is invariant under projective transformations, the final bisector can thus be computed for the canonical case of the unit square and then mapped back using the inverse transformation. $\square$

Given this observation, we will assume henceforth that any four edges case is in the unit square.

Lemma 4.2.3. *When the four edges corresponding to a sector lie on four distinct lines, the bisector is either a hyperbola or a parabola.*

Proof. Let $S(s, t, E_A, E_B, E_C, E_D)$ be an arbitrary sector with four distinct edges. By applying Observation 4.2.2, we may assume without loss of generality that E_A, E_B, E_C, E_D lie on the lines $x = 0, x = 1, y = 0$, and $y = 1$, respectively. The discriminant of equation for our bisector is given by $B^2 - 4AC$. Substituting in for our variables, we get $(1 - k)^2$ as our discriminant, which is nonnegative. $\square$

Next, let us consider the three-edge case. We will refer to the edge that is hit twice by the Hilbert metric in the sector as the *shared edge*, we have the following three cases, depending on whether the shared edge is behind both points with respect to the sector, in front of both (see Figure 4.3), or behind one and in front of the other (see Figure 4.4). As in the four-edge case, we can projectively transform the problem to a canonical form, as given in the following lemma. Define the *unit simplex* in $\mathbb{R}^2$ to be the right triangle with vertices at $(0, 0)$, $(1, 0)$, and $(0, 1)$.

Observation 4.2.4. *Bisectors in the three-edge case can be projectively transformed and solved in the unit simplex.*

Proof. The affine transformation between triangles is well known and inherently projective. We provide it here. A triangle with corners p, q, r can be projectively transformed into the unit simplex with the matrix T^{-1} , where:

$$T = \begin{bmatrix} p_x - r_x & p_y - r_y & 0 \\ q_x - r_x & q_y - r_y & 0 \\ r_x & r_y & 1 \end{bmatrix}.$$

Observe that T^{-1} sends a point P on the left from our original triangle to the unit simplex. □

Lemma 4.2.5. *Given any three-edge sector, where the shared edge is in front or behind of both sites, the bisector between the two sites in the sector is a straight line through the vanishing point of the non-shared edges.*

Proof. We prove this geometrically for both cases $S(s, t, E_A, E_B, E_C, E_D)$ with edges $E_B = E_D$ and $S(s, t, E_A, E_B, E_C, E_D)$ with edges $E_A = E_C$.




Figure 4.3: Illustration for the proof of Lemma 4.2.6 (a bisector in the three-edge case with shared edge in front of both sites).

Consider the case in which the edge is in front of both sites, $S(s, t, E_A, E_B, E_C, E_B)$. Let p be a point in the sector such that $H(s, p) = H(t, p)$, and let O be vanishing point of E_A and E_C . We claim the bisector is a straight line through Op . Let q be a point on the segment Op in the sector. Using the geodesic triangle inequality in the Funk metric we get $F(t, q) = F(t, p) + F(p, q)$ and $F(s, q) = F(s, p) + F(p, q)$ [35]. By definition, $F(s, p) + F(p, s) = F(t, p) + F(p, t)$. Substituting in our previous equations and cancelling $F(p, q)$ we have $F(s, q) + F(p, s) = F(t, q) + F(p, t)$. Using properties of similar triangles and rearranging the equation we can see that $F(s, q) - F(t, q) = \log((|pE_c|/|tE_c|) \cdot (|sE_A|/|pE_A|))$. Note that $|pE_c|/|pE_A| = |qE_c|/|qE_A|$. Substituting in and using the definition of the Funk metric we get $F(s, q) - F(t, q) = F(q, t) - F(q, s)$, yielding $H(s, q) = H(t, q)$. By choosing p to be the lowest point on the bisector in the sector we get our characterization of the bisector as a line.

Note the case where the shared edge is behind both sites follows from the symmetry of the general equation of the bisector. □

Lemma 4.2.6. *Given a three-edge sector in which the shared edge is in front of one site and behind the other, the bisector between the sites can be a conic of any type (ellipse, parabola, or hyperbola)*

depending on their relative positions. Further, when either site is fixed, there exists a line such that the type of conic is determined by the location of the other site with respect to this line (an ellipse if it lies on one side, hyperbola if on the other side, and parabola if lies on the line).

Proof. We prove this algebraically for, without loss of generality, $S(s, t, E_A, E_B, E_C, E_D)$ with edges $E_B = E_C$, and all other edges lying on distinct lines.

By Observation 4.2.4, it suffices to assume that the triangle has been mapped to the unit simplex. Let the edge E_A lie on the line $y = 0$, E_B and E_C lie on the line $x + y - 1 = 0$, and E_D lie on the line $x = 0$. By applying Theorem 4.2.1 with these values, the bisector is given by $x^2 + (2 - k)xy + y^2 - 2x - 2y + 1 = 0$, where $k = \frac{s_x + s_y - 1}{t_x} \frac{t_x + t_y - 1}{s_y}$. The discriminant of the equation of the bisector is $(2 - k)^2 - 4$, which reduces to $k(k - 4)$. By standard results on conics, the type of conic depends on the sign of the discriminant. Note that k is always positive since (s_x, s_y) and (t_x, t_y) are in the unit simplex. Given this, the discriminant changes sign when $k = 4$. Thus, the conic type is given by the sign of the function $(s_x + s_y - 1)(t_x + t_y - 1) - 4s_x t_x$. Observe that when either s or t is fixed, this is a linear equation in coordinates of the other site. Thus, fixing s yields a line, and the type of conic (ellipse, hyperbola, or parabola) is determined by t 's location relative to this line. The elliptical case holds when t is on one side, the hyperbolic case when it is on the other, and the parabolic case when it lies on the line. This applies symmetrically for t . $\square$

Figure 4.4 illustrates four examples of conics as described by Lemma 4.2.6. The site coordinates and discriminant are given in Table 4.1. Note that the ellipse case and hyperbola case happen in general position.

Now that we have evaluated the three-edge case we will evaluate the two-edge case. Note that given a sector $S(s, t, E_A, E_B, E_C, E_D)$ with $E_A = E_D$ and $E_B = E_C$, we can choose a point




Figure 4.4: The bisector conics for the cases given in Table 4.1 visualized.

Table 4.1: The bisector cases illustrated in Figure 4.4 for the three-edge case.

	s	t	Discriminant
Ellipse	(0.2,0.2)	(0.45,0.4)	-3
Hyperbola	(0.2,0.2)	(0.25,0.4)	0.84
Circle	(0.1,0.15)	(0.5,0.3)	-4
Parabola	(0.1,0.1)	(0.2,0.7)	0

p on E_A , q on E_B , and let $E_A \cap E_B = r$ and affinely map the triangle pqr into the unit simplex using Observation 4.2.4.

Lemma 4.2.7. *Given a two edge sector, $S(s, t, E_A, E_B, E_C, E_D)$ with $x = 0$ as $E_A = E_D$, and $y = 0$ as $E_B = E_C$, the bisector between s and t is given by the following linear equation: $y = \sqrt{k}x$ which passes through the vanishing point of E_A and E_B where $k = (s_y/t_x) \cdot (t_y/s_x)$.*

Proof. Plugging in to the general equation for the bisector in a sector give us $-kx^2 + y^2 = 0$ with $k = (s_y/t_x) \cdot (t_y/s_x)$ and solving gives us one viable solution $y = \sqrt{k}x$. It is clear that this goes through the vanishing point of E_A and E_B . $\square$

4.3 Analysis of Hilbert Voronoi Diagrams

The Hilbert metric has some unique features that are not present in the Euclidean metric, In particular, in degenerate cases bisectors in the Hilbert metric can contain two dimensional polygons. Next, we present a necessary and sufficient condition for this occurrence.

Lemma 4.3.1. *The bisector between two sites s and t contains a Hilbert ball of some radius r if and only if both sites lie on a ray through a vertex V at the vanishing point of two edges E_A and E_B of Ω and the sector $S(s, t, E_A, E_B, E_B, E_A)$ exists.*




Figure 4.5: The point p is equidistant from both sites s and t by the invariance of the cross ratio.

Proof. ($\Leftarrow$) The Hilbert distance from s and t to any point p in $S(s, t, E_A, E_B, E_B, E_A)$, from the perspective of both s and t , is equal by the invariance of the cross ratio. Hence, the entire sector in which both s and t are between E_A and E_B is part of the bisector between s and t . By assumption, Ω is a bounded convex polygonal body, so our bisector will contain some fully 2-dimensional region. Note that we can always fit a Euclidean ball in any 2-dimensional region, and by we therefore must be able to fit a Hilbert ball in one as well [35].

($\Rightarrow$) Suppose the bisector between two sites s and t contains a Hilbert ball N . Then there exists a ball N' nested in N so that N' lies entirely within one sector T . Note that the cross ratio between either site and any point in N' must be equivalent. We assert that T is defined by at most two edges of Ω .

Suppose that there were more than two unique edges defining T . Consider a point, p , in N' . Take a pair of edges with respect to one site and consider the ray from the vanishing point of that pair of edges through p . Note that moving p along this line will not change its distance to the site, however, since this sector is defined by more than two edges, it will from the other site. Hence, there can be only two edges of Ω defining T .




Figure 4.6: A discontinuity in the Hilbert metric Voronoi.

By the continuity of the Hilbert Metric, it must be that both sites lie on the same side of the bisector in T . Otherwise, we could move closer to one site and further from the other. Since both sites lie on the same side of their bisector, share two edges with respect to every point in N' , and are the same distance to every point in N' , they must lie on a ray through o . $\square$

Given this, we can see that any dynamic Hilbert Voronoi diagram will contain discontinuities whenever a site passes over a ray from a vanishing point of a pair of edges and another site. (See Figure 4.6 for an example.)

Another property of Hilbert Voronoi diagrams that is distinct from Euclidean Voronoi diagrams is that three sites in general position do not necessarily create a Voronoi vertex. In particular given two site s and t , there exists a space of positive measure such no ball containing a point in that space on its boundary can contain both s and t on its boundary.

Lemma 4.3.2. *There exists a quadrilateral region Z of positive measure between any two sites s and t in a convex polygonal Ω , such that for all $r \in Z$, there is no Hilbert ball whose boundary passes through r , s and t .*

Proof. Let Ω have m edges, denoted $\{E_1, \dots, E_m\}$. Let s and t be two sites in Ω . Take O to be the set of vanishing points on pairs of edges $O = \{O_{1,2}, O_{1,3}, \dots, O_{m-1,m}\}$ when $O_{i,j}$ refers to the

vanishing point of edges E_i and E_j . Let $R_{i,j}$ be the area between rays $O_{i,j}s$ and $O_{i,j}t$ intersected with Ω . Let $R = \{R_{1,2}, R_{1,3}, \dots, R_{m-1,m}\}$.

Let $Z = \bigcap_R R_{i,j}$. By definition, Z is a polygon whose edges pass through both s and t , which always contains st . Note that it is equal to st if and only if s and t lie on a ray through an element of O .

Consider the boundary of Z . We begin by proving that Z is a quadrilateral. Without loss of generality let us consider only the top half of Z above the line st . Let the first two rays leaving s and t and forming part of the boundary of Z be denoted l_s and l_t , respectively. Let l_s and l_t intersect at a point v . If v is not in Z , there must exist some line l that intersects both l_s and l_t and either s or t , and therefore, must be one of the original rays l_s or l_t . This contradicts the fact that v is not in Z . Hence we know Z must be a quadrilateral. Next, we prove that no ball through s and t can intersect the interior of Z on its boundary.

To do this, we prove that any ball containing s and t must contain Z . Consider an arbitrary Hilbert ball N with s and t on its boundary. Note that the boundary of N must be made of segments along rays passing through elements of O . Since s and t are on the boundary N , it must be that if an edge, W of N , intersects Z it would cut through, without loss of generality, the top half of the quadrilateral. Extending W , we see that it must intersect l_s and l_t at exactly one point each. Let W come from a ray passing through $O' \in O$. Consider segments through s and t created by the rays through O' . Note that W must be either above or below $O's$ and $O't$, or otherwise either s or t would not be on the boundary. However, Z is contained in the region between $O's$ and $O't$, so W cannot cut through it, contradiction. Hence we can see that N must contain the interior of Z . $\square$




Figure 4.7: Illustration of region Z.

Chapter 5: Delaunay Triangulations in the Hilbert Metric

5.1 Introduction

As mentioned in the preliminaries to this dissertation, Delaunay triangulation are the graphical dual structure to Voronoi diagrams and are extensively used in their own right. In this chapter we show how to adapt the Euclidean Delaunay triangulation to the Hilbert geometry defined by a convex polygon in the plane. We analyze the geometric properties of the Hilbert Delaunay triangulation, which has some differences with respect to the Euclidean case, including the fact that the triangulation does not necessarily cover the convex hull of the point set. We also introduce the notion of a Hilbert ball at infinity, which is a Hilbert metric ball centered on the boundary of the convex polygon. We present a simple randomized incremental algorithm that computes the Hilbert Delaunay triangulation for a set of n points in the Hilbert geometry defined by a convex m -gon. The algorithm runs in $O(n(\log n + \log^3 m))$ expected time. In addition we introduce the notion of the Hilbert hull of a set of points, which we define to be the region covered by their Hilbert Delaunay triangulation. We present an algorithm for computing the Hilbert hull in time $O(nh \log^2 m)$, where h is the number of points on the hull's boundary.




Figure 5.1: (a) A convex polygon Ω and sites, (b) the Voronoi diagram, and (c) the Delaunay triangulation.

5.2 The Hilbert Delaunay Triangulation

The *Hilbert Delaunay triangulation* of P , denoted $DT(P)$, is defined as the dual of the Hilbert Voronoi diagram, where two sites p and q are connected by an edge if their Voronoi cells are adjacent (see Figure 5.1 (b) and (c)). Throughout, we make the general-position assumption that no four sites of P are Hilbert equidistant from a single point in $\text{int}(\Omega)$. It is easy to see that familiar concepts from Euclidean Delaunay triangulations apply, but using Hilbert balls rather than Euclidean balls. For example, two sites are adjacent in the triangulation if and only if there exists a Hilbert ball whose boundary contains both sites, and whose interior contains no sites. (The center of this ball lies on the Voronoi edge between the sites.) Also, three points define a triangle if and only if there is a Hilbert ball whose boundary contains all three points, but is otherwise empty.

Consider a triangle $\triangle pqr$ in Ω 's interior. A Hilbert ball whose boundary passes through all three points is said to *circumscribe* this triangle. As we shall see in Section 5.4, some triangles admit no circumscribing Hilbert ball, but the following lemma shows that if it does exist, then it is unique. We refer to the boundary of such a ball as a *Hilbert circumcircle*.

Lemma 5.2.1. *There is at most one Hilbert circumcircle for any three non-colinear points in Ω 's interior.*

Proof. Let p, q , and r be the vertices of the triangle, which we assume are not collinear. Suppose for the sake of a contradiction there did exist two circumscribing Hilbert balls B_1 and B_2 . Let c_1 and c_2 denote their respective centers, and let $E = \Omega \setminus (B_1 \cup B_2)$ denote the region exterior to both balls. Since Hilbert balls are convex, p, q , and r all lie on the boundary of E . It follows that for any point w in E , there exist three nonintersecting curves joining w to p, q , and r . Since both balls are empty, points c_1 and c_2 are vertices in the Voronoi diagram of $\{p, q, r\}$. Since Voronoi cells are star-shaped, the six open line segments joining each center to each point do not intersect. Further, these lie within $B_1 \cup B_2$, and so they are disjoint from the curves joining w to these points. Therefore the complete bipartite graph $\{p, q, r\} \times \{c_1, c_2, w\}$ is planar, but this is impossible since it would imply a planar embedding of $K_{3,3}$ (the complete bipartite 3×3 graph). $\square$

Note that the non-collinear assumption is necessary. To see why, let Ω be a axis aligned rectangle, and let ℓ be a horizontal line through the rectangle's center. It is easy to construct two Hilbert balls, one above the line and one below, whose boundaries contain a segment along this line. Placing the three points within this segment would violate the lemma.

Lemma 5.2.2. *Given any discrete set of points P , $\text{DT}(P)$ is a planar (straight-line) graph that spans P .*

Proof. Suppose to the contrary that there existed $a, b, c, d \in P$ such that the edges (a, b) and (c, d) , viewed as open line segments, intersected each other. These edges are witnessed by two empty Hilbert balls, B_1 for (a, b) and B_2 for (c, d) (see Figure 5.2(a)). By convexity, B_1 contains the segment ab , and by emptiness, its interior does not contain c or d . Similarly, B_2 contains the segment cd , and its interior does not contain a or b . and B_2 . It follows the boundaries of B_1 and B_2 must intersect in at least four points. Applying Lemma 5.2.1 to any three of these four intersection

points yields a contradiction.




Figure 5.2: Proof of Lemma 5.2.2.

To prove the $DT(P)$ spans P , suppose to the contrary that this was not so. Consider the two sites $p, q \in P$ that are in different components but closest in terms of their Hilbert distance. Let c denote the center of the smallest Hilbert ball B containing p and q . Because straight lines are geodesics, c lies on the line segment pq , and $H_\Omega(p, q) = H_\Omega(p, c) + H_\Omega(c, q)$ (see Figure 5.2(b)). We assert that this ball can contain no other site in its interior. If there were such a site r , it lies in a different connected component from p or q (possibly both). Let us assume without loss of generality that it is not in p 's connected component. Because it is in the ball's interior, we have $H_\Omega(c, r) < H_\Omega(c, q)$. By the triangle inequality, we obtain

$$H_\Omega(p, r) \leq H_\Omega(p, c) + H_\Omega(c, r) < H_\Omega(p, c) + H_\Omega(c, q) = H_\Omega(p, q),$$

but this contradicts the hypothesis that p and q were the closest sites in different components. $\square$

Another useful property of the Delaunay triangulation, viewed as a graph, is its relationship to other geometric graph structures. The *Hilbert minimum spanning tree* of a point set P , denoted $MST_\Omega(P)$, is the spanning tree over P where edge weights are Hilbert distances. The *Hilbert relative neighborhood graph* for a point set P , denoted $RNG_\Omega(P)$ is a graph over the vertex set P , where two points $p, q \in P$ are joined by an edge if $H_\Omega(p, q) \leq \max(H_\Omega(p, r), H_\Omega(q, r))$, for

all $r \in P \setminus \{p, q\}$. Toussaint proved that in the Euclidean metric, the minimum spanning tree is a subgraph of the relative neighborhood graph, which is a subgraph of the Delaunay graph [43]. The following shows that this holds in the Hilbert metric as well. (Since $\text{MST}_\Omega(P)$ is connected, this provides an alternate proof that $\text{DT}(P)$ spans P .)

Lemma 5.2.3. *Given any discrete set of points P , $\text{MST}_\Omega(P) \subseteq \text{RNG}_\Omega(P) \subseteq \text{DT}_\Omega(P)$.*

Proof. The fact that the minimum spanning tree is contained in the relative neighborhood graph holds in any metric. If two points p and q are not joined by an edge in the RNG, then there is a point r that is closer to both than they are to each other. The edge pq cannot be in the minimum spanning tree, because it is the highest weight edge in the cycle formed by p , q , and r .

To prove that the $\text{RNG}_\Omega(P) \subseteq \text{DT}_\Omega(P)$, consider an edge pq in the Hilbert RNG. Recall $B(p, H_\Omega(p, q))$ denote the Hilbert ball centered at p with radius $H_\Omega(p, q)$ and define $B(q, H_\Omega(q, p))$ analogously. The condition for two points $p, q \in P$ to be connected in $\text{RNG}_\Omega(P)$ is that there is no other point of P lying in the interior of “lune” $B(p, H_\Omega(p, q)) \cap (q, H_\Omega(q, p))$. Consider the ball B centered at the point c that is midway between p and q , as introduced in the proof of Lemma 5.2.2. For any $r \in B$, we showed that $H_\Omega(p, r) < H_\Omega(p, q)$. Symmetrically, we have $H_\Omega(q, r) < H_\Omega(q, p)$, which implies that r lies in the lune. Since the lune must be empty, B must also be empty. This empty ball witnesses the fact that pq is an edge of $\text{DT}_\Omega(P)$. $\square$

5.3 Hilbert Bisectors

In this section we present some utilities for processing Hilbert bisectors, which will be used later in our algorithms. We start by recalling a characterization for when a point lies on the Hilbert bisector from Chapter 3. Recall that three lines are said to be *concurrent* in projective geometry if

they intersect in a common point or are parallel.

Lemma 5.3.1. *Given a convex body Ω and two points $p, q \in \text{int}(\Omega)$, consider any other point $x \in \Omega$. Let p' and p'' denote the endpoints of the chord $\chi(p, x)$ so the points appear in the order $\langle p', p, x, p'' \rangle$. Define q' and q'' analogously for $\chi(q, x)$.*

$\longleftrightarrow \quad \longleftrightarrow \quad \longleftrightarrow$

(i) *If $x \in \text{int}(\Omega)$, then it lies on the (p, q) -bisector if and only if the lines pq , $p'q'$, and $p''q''$ are concurrent (see Figure 5.3(a)).*

(ii) *If $x \in \partial\Omega$ (implying that $x = p'' = q''$), then x is the limit point of the (p, q) -bisector if and only if there is a supporting line through x concurrent with pq and $\overleftrightarrow{p'q'}$ (see Figure 5.3(b)).*




Figure 5.3: Properties of points on the Hilbert bisector.

The following utility lemmas arise as consequences of this characterization. Both involve variants of binary search.

Lemma 5.3.2. *Given an m -sided convex polygon Ω , two points $p, q \in \text{int}(\Omega)$, and any ray emanating from p , the point of intersection between the ray and the (p, q) -bisector (if it exists) can be computed in $O(\log^2 m)$ time.*

Proof. Through binary search on the boundary of Ω , in $O(\log m)$ time we can determine the endpoints p' and p'' of the chord passing through p along the ray (see Figure 5.4(a)). Let ℓ denote this chord, and let V denote the points along the chord where the spokes of q intersect ℓ (see Figure 5.4(b)).

Since V contains $O(m)$ points, and we eliminate a constant fraction with each probe, the overall time is $O(\log^2 m)$, as desired. $\square$

Lemma 5.3.3. *Given an m -sided convex polygon Ω and any two points $p, q \in \text{int}(\Omega)$, the endpoints of the (p, q) -bisector on $\partial\Omega$ can be computed in $O(\log^2 m)$ time.*

Proof. For the sake of illustration, let us assume that the line through p and q , denoted ℓ , is horizontal (see Figure 5.5(a)). In $O(\log m)$ time we can compute the two points, z_0 and z_1 , where ℓ intersects $\partial\Omega$, with z_0 on the left side of Ω and z_1 on the right. These points implicitly split the boundary of Ω into two convex chains, which we call the *upper* and *lower chains*. We will show how to compute the endpoint of the (p, q) -bisector that lies on the upper chain, and the other case is symmetrical.




Figure 5.5: Computing the endpoints of the (p, q) -bisector.

It will be convenient to think of the portion of ℓ that lies outside of Ω as being oriented from right to left, first starting at z_0 , then progressing to left to $-\infty$ and wrapping around to $+\infty$, and finally progressing to z_1 . Observe that as a point x travels from z_0 to z_1 along the upper chain of $\partial\Omega$, the associated supporting line intersects ℓ at a point that moves monotonically along ℓ from z_0 to z_1 (wrapping around when it hits the topmost vertex of Ω).

By Lemma 5.3.1(ii), we seek a point x^* on the upper chain that satisfies the conditions of the lemma (see Figure 5.5(a)). Given any point x on the upper chain, consider the endpoints p' and q'

of the respective chords $\overline{xp}$ and $\overline{xq}$ (see Figure 5.5(b)). Let z denote the point where a supporting line at x intersects ℓ . Let z' denote the point where the line $\overleftrightarrow{p'q'}$ intersects ℓ . In the limit, when $x \rightarrow z_0$, we have $z = z_0$ and $z' = z_1$, and in the other limit, when $x \rightarrow z_1$, we have $z = z_1$ and $z' = z_0$. The key to the search process is that as x moves monotonically along the upper chain from z_0 to z_1 , the point z moves monotonically along ℓ from z_0 to z_1 . On the other hand, the points p' and q' move monotonically along the lower chain from z_1 and z_0 , and hence z' moves in the reverse direction along ℓ . It follows that there is a unique point x on the upper chain where z and z' coincide along ℓ . This will be the desired point x^* .

We will show that x^* can be found through a binary search along the upper chain of Ω , where each probe takes $O(\log m)$ time. We maintain two spokes through p , which delimit the portion along the upper chain where the desired point x^* resides (see Figure 5.5(c)). This partitions the upper and lower chain into two polygonal arcs containing the endpoints of the candidate chords. With each probe, we sample the median vertex x from the arc that contains the larger number of vertices. This will allow us to eliminate at least one quarter of the remaining vertices from consideration.

Depending on the nature of the spoke, the point x may be a vertex of the upper chain or it may lie in the interior of an edge. If it is in the interior of an edge, the supporting line at x is uniquely determined, implying that z is uniquely determined as well. In $O(\log m)$ time, we can determine the points p' and q' , and z' can be computed in $O(1)$ additional time. Depending on where z' lies relative to z in the order defined along ℓ , we eliminate either the region before or after x on the upper chain, thus eliminating a constant fraction of the points from further consideration. On the other hand, if x is a vertex, there are two supporting lines, which define an interval along ℓ . Again, we can determine p' , q' , and z' in $O(\log m)$ time. If z' lies within this interval, then x is

the desired final answer. Otherwise, we eliminate either the region before or after x , depending on which side of the interval z' lies.

Each probe takes $O(\log m)$ time (to determine the points p' and q'). Since there are $O(m)$ spokes through p , and we eliminate a constant fraction with each probe, the overall time is $O(\log^2 m)$, as desired. $\square$

5.4 Hilbert Circumcircles

In the Euclidean plane, any triple of points that are not collinear lie on a unique circle, that is, the boundary of an Euclidean ball. This is not true in the Hilbert geometry, however. Since Delaunay triangulations are based on an empty circumcircle condition, it will be important to characterize when a triangle admits a Hilbert circumcircle and when it does not. In this section we explore the conditions under which such a ball exists. Then we define balls at infinity, which are balls through point

5.4.1 Balls at infinity

We begin by introducing the concept of a Hilbert ball centered at a point on the boundary of Ω . Let x be any point on $\partial\Omega$. Given any point $p \in \text{int}(\Omega)$, we are interested in defining the notion of a Hilbert ball centered at x whose boundary contains p . If we think of the points on the boundary of Ω as being infinitely far from any point in $\text{int}(\Omega)$, this gives rise to the notion of a ball at infinity.

Given $x \in \partial\Omega$ and $p \in \text{int}(\Omega)$, let u be any unit vector directed from x into the interior of Ω (see Figure 5.6(a)). Given any sufficiently small positive δ , let $x_\delta = x + \delta u$ denote the point

at Euclidean distance δ from x along vector u , and let $B(x_\delta, H_\Omega(x_\delta, p))$ denote the Hilbert ball centered at x_δ of radius $H_\Omega(x_\delta, p)$. The following lemma shows that as δ approaches 0, this ball approaches a shape, which we call the *Hilbert ball at infinity* determined by x and u and passing through p , denoted $B_u(x, p)$ (see Figure 5.6(b)).




Figure 5.6: Constructing the Hilbert ball $B_u(x, p)$ at infinity.

Lemma 5.4.1. *As δ approaches 0, $B(x_\delta, H_\Omega(x_\delta, p))$ approaches a convex polygon lying within Ω having x on its boundary.*

A useful utility, which we will apply in Section 5.7, involves computing the largest empty ball centered at any boundary point x with respect to a point set P .

Lemma 5.4.2. *Given a set of n points P in the interior of Ω , and any point $x \in \partial\Omega$, in $O(n \log m)$ time, it is possible to compute a point $p \in P$, such that there is a ball at infinity centered at x whose boundary passes through p , and which contains no points of P in its interior.*

Proof. Let $P = \{p_1, \dots, p_n\}$, and let ℓ be any supporting line passing through x , and fix an arbitrary point $y \in \partial\Omega$ other than x (see Figure 5.7(a)). For each point $p \in P$, shoot a ray from x through p , letting p' denote the point where it hits the boundary. Let v denote the intersection of ℓ and the line $\overleftrightarrow{p'y}$, and let q denote the intersection of line $\overleftrightarrow{p'v}$ with the chord $\overline{xy}$ (see Figure 5.7(b)).




Figure 5.7: Proof of Lemma 5.4.2.

It follows from Lemma 5.3.1 that both p and q lie on the boundary of a ball at infinity centered at x . Applying this construction to each $p_i \in P$ yields a set of points $\{q_1, \dots, q_n\}$ along $\overline{xy}$, which by Lemma 5.2.1 are ordered by inclusion (see Figure 5.7(c)). The desired point $p_i \in P$ is the one whose corresponding point q_i is closest to x along this chord. The running time for each point p_i is $O(\log m)$ to determine where the ray hits the boundary, which yields an overall running time of $O(n \log m)$. $\square$

Balls at infinity are not proper aspects of Hilbert geometry, but they will be convenient for our purposes. Given two points $p, q \in \text{int}(\Omega)$, let x denote the endpoint of the (p, q) -bisector on $\partial\Omega$, oriented so that x lies to the left of the directed line $\overrightarrow{pq}$. (It follows from the star-shapedness of Voronoi cells that the bisector endpoints intersect $\partial\Omega$ on opposite sides of this line.) Let u denote a tangent (if x is a vertex of Ω) vector of the bisector at x . Define $B(p : q) = B_u(x, p)$. Note that this (improper) ball is both centered at and passes through x . In this sense it circumscribes the triangle $\triangle pqx$. Define $B(q : p)$ analogously for the opposite endpoint of this bisector (see Figure 5.8(a) and (b)).

We can now characterize the set of points r that admit a Hilbert circumcircle with respect to two given points p and q . This characterization is based on two regions, called the *overlap* and *outer regions* (see Figure 5.8(c)).

Definition 5.4.3 (Overlap/Outer Regions). Given two points $p, q \in \text{int}(\Omega)$:

Overlap Region: denoted $Z(p, q)$, is $B(p : q) \cap B(q : p)$.

Outer Region: denoted $W(p, q)$, is $\Omega \setminus (B(p : q) \cup B(q : p))$.




Figure 5.8: The overlap region $Z(p, q)$ and the outer region $W(p, q)$.

Lemma 5.4.4. A triangle $\triangle pqr \subseteq \text{int}(\Omega)$ admits a Hilbert circumcircle if and only if $r \notin Z(p, q) \cup W(p, q)$.

As shown in Figure 5.1, the Delaunay triangulation need not cover the convex hull of the set of sites. We refer to the region that is covered as the *Hilbert hull* of the sites. Later, we will present an algorithm for computing the Hilbert hull. The following lemma will be helpful. It establishes a nesting property for the overlap regions.

Lemma 5.4.5. If $r \in Z(p, q)$ then $Z(p, r) \subset Z(p, q)$.

5.5 Computing Circumcircles

A fundamental primitive in the Euclidean Delaunay triangulation algorithm is the so-called *in-circle test* [44]. Given a triangle $\triangle pqr$ and a fourth site s , the test determines whether s lies within the circumscribing Hilbert ball for the triangle (if such a ball exists). In this section, we present an algorithm which given any three sites either computes the Hilbert circumcircle for these

sites, denoted $B(p : q : r)$, or reports that no circumcircle exists. The in-circle test reduces to checking whether $H_\Omega(s, c) < \rho$, which can be done in $O(\log m)$ time.

Lemma 5.5.1. *Given a convex m -sided polygon Ω and triangle $\triangle pqr \subset \text{int}(\Omega)$, in $O(\log^3 m)$ time it is possible to compute $B(p : q : r)$ or to report that no ball exists.*

The remainder of the section is devoted to the proof. The circumscribing ball exists if and only if there is a point equidistant to all three, implying that (p, q) - and (p, r) -bisectors intersect at some point $c \in \text{int}(\Omega)$. The following technical lemma shows that bisectors intersect crosswise.

Lemma 5.5.2. *Given three non-collinear points $p, q, r \in \text{int}(\Omega)$, the (p, q) - and (p, r) -bisectors have endpoints lying on $\partial\Omega$. If they intersect within $\text{int}(\Omega)$, they intersect transversely in a single point (see Figure 5.9).*




Figure 5.9: Circumcircles and bisector intersection.

It follows that the (p, q) - and (p, r) -bisectors subdivide the interior of Ω into either three or four regions (three if they do not intersect, and four if they do). We can determine which is this case by invoking Lemma 5.3.3 to compute the endpoints of these bisectors in $O(\log^2 m)$ time. If they alternate between (p, q) and (p, r) along the boundary of Ω , then the bisectors intersect, and otherwise they do not. In the latter case, there is no circumcircle for p, q , and r , and hence, no possibility of violating the circumcircle condition. (Note that this effectively provides an $O(\log^2 m)$ test for Lemma 5.4.4.) Henceforth, we concentrate on the former case.

Let us assume, without loss of generality, that $\triangle pqr$ is oriented counterclockwise. Let v_q denote the endpoint of the (p, q) -bisector that lies to the right of the oriented line $\overrightarrow{pq}$. Let v_r denote the endpoint of the (p, r) -bisector that lies to the left of the oriented line $\overrightarrow{pr}$ (see Figure 5.10(a)).




Figure 5.10: Computing the center of $\triangle pqr$.

Because Voronoi cells are star-shaped, it follows that the vector from p to the desired circumcenter c lies in the counterclockwise angular interval from $\overrightarrow{pv_q}$ to $\overrightarrow{pv_r}$ (see Figure 5.10(b)). The key to the search is the following observation. In the angular region from $\overrightarrow{pv_q}$ to $\overrightarrow{pc}$, any ray shot from p intersects the (p, q) -bisector before hitting the (p, r) -bisector (if it hits the (p, r) -bisector at all). On the other hand, in the angular region from $\overrightarrow{pc}$ to $\overrightarrow{pv_r}$, any ray shot from p intersects the (p, r) -bisector before hitting the (p, q) -bisector (if it hits the (p, q) -bisector at all). We say that the former type of ray is *early* and the latter type is *late*.

Let v_- and v_+ denote the points on $\partial\Omega$ that bound the current search interval about p (see Figure 5.10(c)). We will maintain the invariant that the ray $\overrightarrow{pv_-}$ is early and $\overrightarrow{pv_+}$ is late. Initially, $v_- = v_q$ and $v_+ = v_r$. Let v'_- and v'_+ denote the opposite endpoints of the chords χpv_- and χpv_+ . Consider the portion of the boundary of Ω that lies counterclockwise from v_- and v_+ . If the angle $\angle v_-pv_+$ is smaller than π , also consider the portion of the boundary of Ω that lies counterclockwise from v'_- and v'_+ . With each probe, we sample the median vertex v from whichever of these two boundary portions that contains the larger number of vertices. If v comes from the interval $[v_-, v_+]$,

we probe along the ray $\overrightarrow{pv}$, and if it comes from the complementary interval, $[v'_-, v'_+]$, we shoot the ray from p in the opposite direction from v . We then apply Lemma 5.3.2 twice to determine in $O(\log^2 m)$ time where this ray hits the (p, q) - and (p, r) -bisectors (if at all). Based on the results, we classify this ray as being early or late and recurse on the appropriate angular subinterval. When the search terminates, we have determined a pair of consecutive spokes about p that contain c .

Because each probe eliminates at least half of the vertices from the larger of the two boundary portions, it follows that at $O(\log m)$ probes, we have located c to within a single pair of consecutive spokes around p . Since each probe takes $O(\log^2 m)$ time, the entire search takes $O(\log^3 m)$ time.

We repeat this process again for r and q . The result is three double wedges defined by consecutive spokes, one about each site. It follows from our earlier remarks from Chapter 3 that within the intersection of these regions, the bisectors are simple conics. We can compute these conics in $O(1)$ time [32] and determine their intersection point, thus yielding the desired point c . The radius ρ of the ball can also be computed in $O(1)$ time.

5.6 Building the Triangulation

In this section we present our main result, a randomized incremental algorithm for constructing the Delaunay triangulation $DT(P)$ for a set of n sites P in the interior of an m -sided convex polygon Ω . Our algorithm is loosely based on a well-known randomized incremental algorithm for the Euclidean case [38, 44].

5.6.1 Orienting and Augmenting the Triangulation

In this section we introduce some representational conventions for the sake of our algorithm. First, we will orient the elements of the triangulation. Given $p, q \in \text{int}(\Omega)$ define the *endpoint* of the (p, q) -bisector to be the endpoint that lies to the left of the directed line $\overrightarrow{pq}$. (It is worth repeating that it follows from the star-shapedness of Voronoi cells that the bisector endpoints intersect $\partial\Omega$ on opposite sides of this line.) The opposite endpoint will be referred to as the (q, p) -bisector endpoint. When referring to a triangle $\triangle pqr$, we will assume that the vertices are given in counterclockwise order. Also, edges in the triangulation are assumed to be directed, so we can unambiguously reference the left and right sides of a directed edge pq .

As mentioned earlier, the triangulation need not cover the convex hull of the set of sites (see Figure 5.11(a)). For the sake of construction, it will be convenient to augment the triangulation with additional elements, so that all of Ω is covered. First, we add elements to include the endpoints of the Voronoi bisectors on $\partial\Omega$. For each edge pq such that the external face of the triangulation lies to its left, the endpoint of the (p, q) -bisector intersects the boundary of Ω . Letting x denote this boundary point, we add a new triangle $\triangle pqx$, called a *tooth*. (See the green shaded triangles in Figure 5.11(b).)




Figure 5.11: The augmented triangulation.

Finally, the portion of Ω that lies outside of all the standard triangles and teeth consists of a collection of regions, called *gaps*, each of which involves a single site, the sides of two teeth, and a convex polygonal chain along $\partial\Omega$. While these shapes are not triangles, they are defined by three points, and we will abuse the notation $\triangle pxy$ to refer to the gap defined by the site p and boundary points x and y (see Figure 5.11(b)).

Even when sites are in general position, multiple teeth can share the same boundary vertex. For example, in Figure 5.11(b) three teeth meet at the same boundary vertex v . In our augmented representation, it will be convenient to treat these as three separate teeth, meeting at three distinct (co-located) vertices, separated by two degenerate (zero-width) gaps. This allows us to conceptualize the region outside of the standard triangulation as consisting of an alternating sequence of teeth and gaps. The following lemma summarizes a few useful facts about the augmented representation.

Lemma 5.6.1. *Given a set of n point sites in the interior of a convex polygon Ω :*

1. *The augmented Delaunay triangulation has complexity $O(n)$.*
2. *The region covered by standard triangles is connected.*
3. *Each standard triangle and each tooth satisfies the empty circumcircle property.*
4. *The region outside the standard triangles consists of an alternating sequence of teeth and gaps.*
5. *For any $p \in \Omega$, membership in any given triangle, tooth, or gap can be determined in $O(1)$ time.*

Proof. Claim 1. holds because the standard triangulation is clearly $O(n)$ by standard results on triangulations (based on Euler's formula). The number of teeth is clearly $O(n)$, since each is generated by a standard edge. Claim 2. holds via the connectivity of the triangulation which

was proved in Lemma 5.2.2. Claim 3. and 4 follow through the addition of degenerate gaps. By construction the teeth and gaps alternate. Note membership in a standard triangle or tooth can be computed in $O(1)$ time. Claim 5. holds since each gap is bounded by two edges and the boundary of Ω . Hence membership for any $p \in \Omega$ can be determined in $O(1)$ time. $\square$

5.6.2 Local and Global Delaunay

Given a set P of point sites, we say that an augmented triangulation $\mathcal{T}(P)$ is *globally Delaunay* (or simply *Delaunay*) if for each standard triangle and each tooth $\triangle pqr$, the circumscribing ball, denoted $B(p:q:r)$, exists and has no site within its interior. The incremental Euclidean Delaunay triangulation algorithm employs a local condition, which only checks this for neighboring triangles [38, 44]. Given a directed edge pq of the triangulation that is incident to two triangles or to a triangle and tooth, let a and b be the vertices of the incident triangles lying to the left and right of the pq , respectively. We say that $\mathcal{T}(P)$ is *locally Delaunay* if $a \notin \text{int}(B(q:p:b))$ and $b \notin \text{int}(B(p:q:a))$ for all such edges. The following lemma shows that it suffices to certify the Delaunay properties locally.

Lemma 5.6.2. *Given a set P of point sites, an augmented triangulation $\mathcal{T}(P)$ is globally Delaunay if and only if it is locally Delaunay.*

Proof. Our proof follows the same structure as that given in [38] for the Euclidean case. Suppose to the contrary that $\mathcal{T}(P)$ is locally but not globally Delaunay. That is, there exist sites $a, b, c, d \in P$ such that $\mathcal{T}(P)$ has a triangle or tooth $\triangle abc$ such that the circumscribing ball $B(a:b:c)$ contains d in its interior (see Figure 5.12(a)). Let ab be the edge such that c and d lie on opposite sides of $\overline{ab}$. Among all triangle-site pairs $(\triangle abc, d)$, let this be the one that maximizes $\angle adb$.




Figure 5.12: Proof of Lemma 5.6.2.

To complete the proof, we consider two cases. First, if $\triangle abc$ is a standard triangle or if it is a tooth and both a and b are sites, consider the adjacent triangle $\triangle bae$ along edge ab . Since $\mathcal{T}(P)$ is locally Delaunay, e does not lie within the interior of $B(a:b:c)$. By Lemma 5.2.1, the circumscribing ball $B(b:a:e)$ completely contains the portion of $B(a:b:c)$ that lies on the same side of the chord $\overline{ab}$ as e . (If it failed to do so, there would be an additional point of intersection of the boundaries of these balls, other than a and b , violating the lemma.) This implies that $d \in B(b:a:e)$. Now, let be denote the edge of $\triangle bae$ such that a and d lie on the opposite sides of the chord $\overline{be}$. It follows that $\angle edb > \angle adb$, contradicting the definition of the pair $(\triangle abc, d)$.

In the second case, $\triangle abc$ is a tooth, and either a or b is a boundary vertex, implying that there is a gap on the other side of ab . In this case, relabel the vertices so that a is a site and b is on the boundary. Let $\triangle b'ae$ denote the tooth adjacent to the opposite side of this gap (see Figure 5.12(b)). Now, apply the above, using $\triangle b'ae$ in place of $\triangle bae$. The key property still holds, namely that circumscribing ball $B(b':a:e)$ completely contains the portion of $B(a:b:c)$ that lies on the same side of the chord $\overline{ab}$ as e . $\square$

5.6.3 Incremental Construction

The algorithm operates by first randomly permuting the sites. It begins by inserting two arbitrary sites a and b and generating the resulting augmented triangulation. In $O(\log^2 m)$ time,

we can compute the endpoints of the (a, b) -bisector. We add the edge ab and create two teeth by connecting a and b to the bisector endpoints (see Figure 5.13(a)). We then insert the remaining points one by one, updating the triangulation incrementally after each insertion (see Algorithm 1). Later, we will discuss how to determine which element each new site lies in, but for now, let us focus on how the triangulation is updated.

Algorithm 1 Constructs the Hilbert Delaunay triangulation of a point set P

```

procedure DELAUNAY( $P$ )
     $\mathcal{T} \leftarrow$  empty triangulation
    Randomly permute  $P$ 
     $a, b \leftarrow$  any two points of  $P$ 
     $x, y \leftarrow$  endpoints of the  $(a, b)$ -bisector
    Add edges  $ab, ax, ay, bx, by$  to  $\mathcal{T}$ 
    for all  $p \in P \setminus \{a, b\}$  do
        INSERT( $p, \mathcal{T}$ )
    end for
    return  $\mathcal{T}$ 
end procedure

```




Figure 5.13: (a) Initialization and insertion into a (b) standard triangle, (c) tooth, (d) gap.

When we insert a point, there are three possible cases, depending on the type of element that contains the point, a standard triangle (three sites), a tooth (two sites), or a gap (one site). The case for standard triangles is the same as for Euclidean Delaunay triangulations [44], involving connecting the new site to the triangle's vertices (see Figure 5.13(b)). Insertion into a tooth involves removing the boundary vertex, connecting the new site to the two existing site vertices, and creating two new teeth based on the bisectors to these sites (see Figure 5.13(c)). Finally, insertion into a

gap involves connecting the new site to the existing site vertex, and creating two new teeth based on the bisectors to this site (see Figure 5.13(d)).

Algorithm 2 Site insertion

```

procedure INSERT( $p, \mathcal{T}$ )                                ▷ Insert a new site  $p$  into triangulation  $\mathcal{T}$ 
   $\triangle abc \leftarrow$  the triangle of  $\mathcal{T}$  containing  $p$ 
  if  $\triangle abc$  is a standard triangle then                  ▷ See Figure 5.13(b)
    Add edges  $ap, bp, cp$  to  $\mathcal{T}$ 
    FLIPEDGE( $ab, p, \mathcal{T}$ ); FLIPEDGE( $bc, p, \mathcal{T}$ ); FLIPEDGE( $ca, p, \mathcal{T}$ )
  else if  $\triangle abc$  is a tooth then                        ▷ See Figure 5.13(c)
    Let  $a$  and  $b$  be sites, and  $c$  be on boundary
     $x, y \leftarrow$  endpoints of the  $(a, p)$ - and  $(p, b)$ -bisectors, respectively
    Remove  $c$  and edges  $ca$  and  $cb$  from  $\mathcal{T}$ 
    Add edges  $ap, bp, ax, px, by, py$  to  $\mathcal{T}$ 
    FLIPEDGE( $ab, p, \mathcal{T}$ );
    FIXTOOTH( $\triangle apx, p, \mathcal{T}$ ); FIXTOOTH( $\triangle pby, p, \mathcal{T}$ )
  else                                                       ▷  $\triangle abc$  is a gap; see Figure 5.13(d)
    Let  $a$  be the site, and let  $b$  and  $c$  be on the boundary
     $x, y \leftarrow$  endpoints of  $(a, p)$ - and  $(p, a)$ -bisectors, respectively
    Add edges  $ap, ax, px, ay, py$  to  $\mathcal{T}$ 
    FIXTOOTH( $\triangle apx, p, \mathcal{T}$ ); FIXTOOTH( $\triangle pay, p, \mathcal{T}$ )
  end if
end procedure

```

On return from INSERT, $\mathcal{T}$ is a topologically valid augmented triangulation, but it may fail to satisfy the Delaunay empty circumcircle conditions, and may even fail to be geometrically valid. The procedures FLIPEDGE and FIXTOOTH repair any potential violations of the local Delaunay conditions.

The procedure FLIPEDGE is applied to all newly generated standard triangles $\triangle pab$. It is given the directed edge ab of the triangle, such that p lies to the left of this edge. It accesses the triangle $\triangle abc$ lying to the edge's right. This may be a standard triangle or a tooth. In either case, it has an associated circumcircle, which we assume has already been computed.¹ We test whether p

¹In the Euclidean case, the in-circle test may be run from either $\triangle abc$ or $\triangle pab$, but this is not true for the Hilbert geometry. In light of Lemma 5.4.4, we do not know whether the $\triangle pab$ has a circumcircle, so we run the test from $\triangle abc$.

encroaches on this circumcircle, and if so, we remove the edge ab . If $\triangle abc$ is a standard triangle, then we complete the edge-flip by adding edge pc , and then continue by checking the edges ac and cb (see Figure 5.14(a)). Otherwise, we remove vertex c , and compute the endpoints x and y of the (p, a) - and (b, p) -bisectors, respectively. We create two new teeth, $\triangle pax$ and $\triangle bpy$ (see Figure 5.14(b)). This creates a new (possibly degenerate) gap $\triangle pxy$. Later, we will show (see Lemma 5.6.3) that no further updates are needed. The algorithm is presented in Algorithm 3.




Figure 5.14: (a) Standard-triangle edge flip, (b) tooth edge flip, and (c) fixing a tooth.

Algorithm 3 In-circle test and edge flip.

```

procedure FLIPEDGE( $ab, p, \mathcal{T}$ )           ▷ In-circle test. New site  $p$  lies to the left of edge  $ab$ .
 $c \leftarrow$  vertex of triangle to right of  $ab$  in  $\mathcal{T}$ 
if  $p \in B(a : c : b)$  then                 ▷  $p$  fails the in-circle test for  $\triangle acb$ 
    Remove edge  $ab$  from  $\mathcal{T}$ 
    if  $\triangle acb$  is a standard triangle then   ▷ See Figure 5.14(a)
        Add edge  $pc$  to  $\mathcal{T}$ 
        FLIPEDGE( $ac, p, \mathcal{T}$ ); FLIPEDGE( $cb, p, \mathcal{T}$ )   ▷ Create  $\triangle pac$  and  $\triangle bpc$ 
    else                                       ▷  $\triangle acb$  is a tooth. See Figure 5.14(b)
         $x, y \leftarrow$  endpoints of  $(p, a)$ - and  $(b, p)$ -bisectors, respectively
        Remove  $c$  and edges  $cb$  and  $ca$  from  $\mathcal{T}$ 
        Add edges  $ax, px, by, py$  to  $\mathcal{T}$            ▷ Create teeth  $\triangle pax$  and  $\triangle bpy$ 
    end if
end if
end procedure

```

The “else” clause creates two teeth $\triangle pax$ and $\triangle bpy$. The following lemma shows that there is no need to apply FIXTOOTH to them.

Lemma 5.6.3. *On return from FLIPEDGE the teeth $\triangle pax$ and $\triangle bpy$ generated in the “else” clause are both valid.*

Proof. Consider the two teeth lying immediately clockwise and counterclockwise from c along the boundary of Ω . Let x' and y' denote the respective boundary vertices of these two teeth. We will show that x lies within the portion of the boundary of Ω running clockwise from c to x' . It will follow from a symmetrical argument that y lies in the counterclockwise portion of the boundary from c to y' .

By the preconditions of FLIPEGE, we know that p lies to the left of edge ab , and since we have executed the “else” clause, we know that p encroaches on the circumscribing ball $B(a : c : b)$. Let $\triangle adx'$ denote the tooth lying clockwise from c prior to the insertion of p . If x lies between the c and x' , then the circumscribing balls for these triangles can be arranged as shown in Figure 5.15(a).




Figure 5.15: Proof of Lemma 5.6.3.

Suppose, to the contrary, that x lies on the other side of x' (see Figure 5.15(b)), observe that the circumcircle $B(p : a : x)$ must cross through the circumcircle $B(a : d : x')$. It follows that, in addition to a and d , these two Hilbert circumcircles intersect in at least one other point z . However, this means that there are two different Hilbert circumcircles passing through the three points a , d , and z , which cannot happen according to Lemma 5.2.1. $\square$

Finally, we present FIXTOOTH. To understand the issue involved, consider Figure 5.13(d). In the figure, the newly created teeth $\triangle pax$ and $\triangle apy$ both lie entirely inside the existing gap. But, this need not generally be the case, and the newly created boundary vertex may lie outside the current gap, resulting in two or more teeth that overlap each other (see, e.g., Figure 5.14(c)). The

procedure **FIXTOOTH** is given a tooth $\triangle abx$, where a and b are sites, and x is on Ω 's boundary. Whenever it is invoked the new site p is either a or b . Recalling our assumption that teeth and gaps alternate around the boundary of Ω , we check the teeth that are expected to be adjacent to the current tooth on the clockwise and counterclockwise sides. (Only the clockwise case is presented in the algorithm, but the other case is symmetrical, with a and b swapped.)

Let $\triangle bcy$ denote the tooth immediately clockwise around the boundary. We test whether these triangles overlap (see Figure 5.14(c)). If so, we know that the (a, b) -bisector (which ends at x) and the (b, c) -bisector (which ends at y) must intersect. This intersection point is the center of a Hilbert ball circumscribing $\triangle abc$. We replace the two teeth $\triangle abx$ and $\triangle bcy$ with the standard triangle $\triangle abc$ and the tooth $\triangle acz$, where z is the endpoint of the (a, c) -bisector. If $p = a$, then we invoke **FLIPEDGE** on the opposite edge bc from p , and we invoke **FIXTOOTH** on the newly created tooth. When the algorithm terminates, all the newly generated elements have been locally validated. The algorithm is presented in Algorithm 4.

Algorithm 4 Check for and fix overlapping teeth.

```

procedure FIXTOOTH( $\triangle abx, p, \mathcal{T}$ )                                ▷ Fix tooth  $\triangle abx$  where  $p = a$  or  $p = b$ 
  Let  $\triangle bcy$  be the tooth clockwise adjacent to  $\triangle abx$ 
  if  $\triangle abx$  overlaps  $\triangle bcy$  then                                ▷ See Figure 5.14(c)
     $z \leftarrow$  endpoint of the  $(a, c)$ -bisector
    Remove  $x$  and  $y$  and edges  $ax, bx, by$ , and  $cy$  from  $\mathcal{T}$ 
    Add edges  $ac, az$ , and  $cz$  to  $\mathcal{T}$                                 ▷ Create triangle  $\triangle abc$  and tooth  $\triangle acz$ 
    if  $p = a$  then
      FLIPEDGE( $bc, p, \mathcal{T}$ )                                ▷ Check edge flip with triangle opposite  $bc$ 
      FIXTOOTH( $\triangle acz, p, \mathcal{T}$ )                                ▷ Check for further overlaps
    end if
  else
    Repeat the above for the triangle counterclockwise from  $\triangle abx$  swapping  $a \leftrightarrow b$ 
  end if
end procedure

```

Based on our earlier remarks, it follows that our algorithm correctly inserts a site into the

augmented triangulation.

Lemma 5.6.4. *Given an augmented triangulation $\mathcal{T}(P)$ for a set of sites P , the procedure INSERT correctly inserts a new site p , resulting in the augmented Delaunay triangulation for $P \cup \{p\}$.*

The final issue is the algorithm's expected running time. Our analysis follows directly from the analysis of the randomized incremental algorithm for the Euclidean Delaunay triangulation. We can determine which triangle contains each newly inserted point in amortized $O(\log n)$ expected time either by building a history-based point location data structure (as with Guibas, Knuth, and Sharir [44]) or by bucketing sites (as with de Berg *et al.* [38]). The analysis in the Euclidean case is based on a small number of key facts, which apply in our context as well. First, sites are inserted in random order. Second, the conflict set for any triangle or tooth consists of the points lying in the triangle's circumcircle. Both of these clearly hold in the Hilbert setting. Third, the number of structural updates induced by the insertion of site p is proportional to the degree of p following the insertion. This holds for our algorithm, because each modification to the triangulation induced by p 's insertion results in a new edge being added to p (and these edges are not deleted until future insertions). Fourth, the structure is invariant to the insertion order. Finally, the triangulation graph is planar, and hence it has constant average degree. The principal difference is that the in-circle test involves computing a Hilbert circumcircle, which by Lemma 5.5.1 can be performed in $O(\log^3 m)$ time. These additional factors of $O(\log n)$ and $O(\log^3 m)$ are performed a constant number of times in expectation, for each of the n insertions. This implies our main result.

Theorem 5.6.5. *Given a set of n points in the Hilbert geometry defined by a convex m -gon Ω , it is possible to construct the augmented Delaunay triangulation in randomized expected time $O(n(\log n + \log^3 m))$.*

5.7 The Hilbert Hull

As observed earlier, the triangles of the Delaunay triangulation of P do not necessarily cover the convex hull of P . The region covered by these triangles is called the *Hilbert hull* (the blue region of Figure 5.11). In this section, we present a simple algorithm for computing this hull for a set of n points in the Hilbert distance defined by a convex m -gon Ω as an alternative to deriving it from the Delaunay triangulation of the point-set. Our approach is roughly based on the Jarvis march algorithm [45] for computing convex hulls.

The standard Jarvis march maintains two consecutive sites on the hull, and it iteratively computes the next point in $O(n)$ time using an angular ordering induced by these two points. Our algorithm will instead maintain a current site p on the hull, and a boundary point x , such that there is an empty ball at infinity centered at x whose boundary passes through p . It uses x and p to induce an angular order to select the next point.

To start the process off, we need to identify a pair (p_0, x_0) , where $p_0 \in P$, $x_0 \in \partial\Omega$, and p_0 lies on the boundary of an empty ball centered at x_0 . By Lemma 5.4.2, such a pair can be computed in $O(n \log m)$ time. Assuming that the algorithm has generated an initial sequence of points on the hull, arriving at a pair (p_i, x_i) satisfying the above invariant, we compute the next pair as follows. First, we take r to be the point of P that minimizes the counterclockwise angle $\angle x_i p_i r$. Such a point has not already been added to the hull.

We then enumerate the points of $P \setminus \{p_i, r\}$ to see whether any lie in the overlap region $Z(p, r)$. Whenever we encounter such a point, we set r to this point and continue the process. After considering all the points of P , it follows from nesting properties of overlap regions (see Lemma 5.4.5) that $Z(p, r)$ is empty, and hence the (r, p) -bisector extends to the boundary of Ω in

the Voronoi diagram of P . Thus, we take $p_{i+1} \leftarrow r$, and x_{i+1} is the bisector endpoint. From the remarks made after Lemma 5.5.2, we can test membership in the $Z(p, r)$ in $O(\log^2 m)$ time, and by Lemma 5.3.3, we can compute the endpoint on the bisector in $O(\log^2 m)$ time as well. Since we consider at most n points, it takes a total of $O(n \log^2 m)$ time to generate one more point on the Hilbert hull. In summary we have that the overall running time is $O(nh \log^2 m)$.

Lemma 5.7.1. *Given a set of n points in the Hilbert geometry defined by a convex m -gon Ω , it is possible to construct the Hilbert hull in time $O(nh \log^2 m)$, where h is the number of points on the hull's boundary.*

Chapter 6: Voronoi Diagrams in the Funk Metric

6.1 Introduction

In this chapter, we return to Voronoi diagrams. As mentioned in the preliminaries, it is of use to consider the Hilbert and Funk metric in convex cones. Here we study Voronoi diagrams under the forward and reverse Funk metrics in polygonal and elliptical cones. We establish several key geometric properties: (1) bisectors consist of a set of rays emanating from the apex of the cone, and (2) Voronoi diagrams in the d -dimensional forward (or reverse) Funk metrics are equivalent to additively-weighted Voronoi diagrams in the $(d - 1)$ -dimensional forward (or reverse) Funk metric on bounded cross sections of the cone. Leveraging this, we provide an $O(n^{\lceil \frac{d-1}{2} \rceil + 1})$ -time algorithm for creating these diagrams in d -dimensional elliptical cones using a transformation to and from Apollonius diagrams, and an $O(mn \log(n))$ algorithm for 3-dimensional polygonal cones with m facets via a reduction to abstract Voronoi diagrams. We also provide a complete characterization of when three sites have a circumcenter in 3-dimensional cones. This and a concurrent work [46] are the first algorithmic studies of the Funk metrics.

6.2 Computing Funk Voronoi Diagrams in Cones

The (s, s') –bisector is the set of points equidistant to both sites s_1 and s_2 . Since the Funk metrics are asymmetric, we define the following conventions:

1. The *forward Funk* Voronoi cell of s is $\{q \in C \mid F_C(s, q) \leq F_C(s', q) \text{ for all } s' \in S\}$
2. The *reverse Funk* Voronoi cell of s is $\{q \in C \mid F_C^r(s, q) \leq F_C^r(s', q) \text{ for all } s' \in S\}$

6.2.1 Empty Voronoi Cells

We begin by studying Voronoi diagrams in the Funk metrics in arbitrary convex cones. Throughout we make the general position assumption that no site lies on the boundary of a cone centered at another site. We first characterize when a site has an empty Voronoi cell. In the next section, we show that forward and reverse Funk Voronoi cells are star-shaped. We assume that the cone is convex with apex at the origin.

Lemma 6.2.1. *Given sites $a, b \in \text{int } C \subset \mathbb{R}^d$, if $b \in \text{int } C_a$ then for all $c \in \text{int } C$, $F_C(a, c) < F_C(b, c)$ (see Figure 2.5(a)). In this case we say that a **dominates** b in the Funk metric.*

Proof. Recall from Definition 2.5.2,

$$F_C(a, c) = \ln \inf\{\beta \in \mathbb{R} \mid \beta c \in C_a\} \text{ and } F_C(b, c) = \ln \inf\{\beta \in \mathbb{R} \mid \beta c \in C_b\}.$$

Since $b \in \text{int } C_a$, we have $C_b \subsetneq C_a$ so every β satisfying $\beta c \in C_b$ also satisfies $\beta c \in C_a$.

Therefore, the infimum over the larger set $\{\beta : \beta c \in C_a\}$ is smaller: $F_C(a, c) < F_C(b, c)$. $\square$

From this lemma, we obtain the following immediate consequence.

Corollary 6.2.2. *Let S be a set of sites in C . If $a \leq_C b$ for some $a, b \in S$, then b has an empty Voronoi cell in the Voronoi diagram of S in the forward Funk.*

Lemma 6.2.3. *Given two sites $a, b \in \text{int } C \subset \mathbb{R}^d$, if $b \in \text{int } C_a^r$ then for all $c \in \text{int } C$, $F_C^r(a, c) < F_C^r(b, c)$. In this case we say that a **dominates** b in the reverse Funk metric.*

Proof. Since $b \in \text{int } C_a^r$, $a \in C_b$. Hence, $a \geq_C b$. Now consider a third point $c \in \text{int } C$. Let η be the infimum value such that $c \leq_C \eta b$, that is, $\eta b - c \in C$. Algebraically, we can see that

$$\eta a - c = \eta a + \eta b - \eta b - c = \eta(a - b) + (\eta b - c).$$

Since both $\eta(a - b)$ (as $a - b \in \text{int } C$) and $\eta b - c$ are in C we know $\eta a - c \in \text{int } C$. Hence $\eta a \in C_c$. So: $\inf\{\beta : c \leq_C \beta a\} < \inf\{\beta : c \leq_C \beta b\}$. By taking logarithms: $F_C(c, a) < F_C(c, b)$, and $F_C^r(a, c) < F_C^r(b, c)$. $\square$

Corollary 6.2.4. *Let S be a set of sites in C . If $b \leq_C a$ for some $a, b \in S$, then b has an empty Voronoi cell in the Voronoi diagram of S in the reverse Funk.*

6.2.2 Voronoi Cells

In this section, and future sections, we will denote the Voronoi cell of a site p as $V(p)$. We prove that Voronoi cells are star-shaped. Recall that a set is star-shaped if there exists a point p in it such that for all q in the set, the line segment from p to q is contained in the set.

Lemma 6.2.5. *The Voronoi cell of a weighted site (p, w_p) in the forward and reverse Funk metrics in a cone C is star-shaped with respect to p (see Figure 6.1 (b)).*

Proof. Let us consider the forward Funk. Suppose towards a contradiction that there exists a site (p, w_p) and points $x, y \in \text{int } C$ such that $x \in V(p)$, $y \notin V(p)$, and y lies on the line segment p to x . Since x and y lie on the same ray from p , with y between p and x , by the geodesic property of straight lines (Lemma 2.5.6):

$$F_C(p, x) = F_C(p, y) + F_C(y, x).$$

Let (q, w_q) be the site closest to y . Since $y \notin V(p)$, we have $F_C(q, y) + w_q < F_C(p, y) + w_p$. Since $x \in V(p)$, we have $F_C(p, x) + w_p \leq F_C(q, x) + w_q$. Combining these inequalities:

$$F_C(q, x) + w_q \geq F_C(p, x) + w_p = F_C(p, y) + F_C(y, x) + w_p > F_C(q, y) + w_q + F_C(y, x).$$

This contradicts the triangle inequality for the Funk metric. Therefore, Voronoi cells are star-shaped with respect to their sites. This applies equally across bounded cross sections. The proof is analogous for the reverse Funk. □




Figure 6.1: (a) Voronoi cells in the reverse Funk metric, and (b) in the forward Funk metric.

6.2.3 Voronoi Algorithms

6.2.3.1 Bisectors as Rays Through the Origin

We now proceed to discuss bisectors. We start with the additive properties of the Funk metrics and use them to show that bisectors in the Funk metrics are made out of rays through the apex of the cone (the origin).

Lemma 6.2.6. *Let $p, q \in \text{int } C$ and $\lambda > 0$. Then the forward and reverse Funk metrics satisfy the following additive decomposition along rays from the apex of the cone:*

1. $F_C(p, q) = F_C(p, \lambda p) + F_C(\lambda p, q) = -\ln(\lambda) + F_C(\lambda p, q)$
2. $F_C^r(p, q) = F_C^r(p, \lambda p) + F_C^r(\lambda p, q) = \ln(\lambda) + F_C^r(\lambda p, q).$

Proof. Consider the 2D plane through p, q , and the origin. Let R be the boundary ray of C in this plane such that the forward Funk distance from p to q is computed using R (via the similar triangles formula). Since p and λp lie on the same ray from the origin, all three points $p, \lambda p$, and q lie in this same plane, and we can compute all distances using the same boundary ray R (by similar triangles the distance between p and λp could be calculated with respect to any ray through the apex). We have:

$$\begin{aligned} F_C(p, \lambda p) + F_C(\lambda p, q) &= \ln \frac{d_2(p, R)}{d_2(\lambda p, R)} + \ln \frac{d_2(\lambda p, R)}{d_2(q, R)} = \ln \frac{d_2(p, R)}{d_2(\lambda p, R)} \frac{d_2(\lambda p, R)}{d_2(q, R)} \\ &= \ln \frac{d_2(p, R)}{d_2(q, R)} = F_C(p, q). \end{aligned}$$

Note that $\ln \frac{d_2(p, R)}{d_2(\lambda p, R)} = -\ln(\lambda)$. For the reverse Funk we recall that $F_C^r(a, b) = F_C(b, a)$. Using a

similar process as before:

$$\begin{aligned} F_C^r(p, q) &= F_C(q, p) = F_C(q, \lambda p) + F_C(\lambda p, p) \\ &= F_C^r(\lambda p, q) + F_C^r(p, \lambda p) = F_C^r(p, \lambda p) + F_C^r(\lambda p, q). \end{aligned}$$

The rest follows from: $F_C^r(p, \lambda p) = F_C(\lambda p, p) = \ln(\lambda)$. □

Adding these provides an alternative proof of the Hilbert metric's projective invariance. Next we show that bisectors in the Funk metrics are made of rays through the origin.

Theorem 6.2.7. *The bisector between two sites $a, b \in \text{int } C$ in the forward and reverse Funk metrics consists of a set of rays emanating from the origin.*

Proof. To show that bisectors consist of a set of rays emanating from the origin, we will prove that if $F_C(a, p) = F_C(b, p)$, then $F_C(a, \lambda p) = F_C(b, \lambda p)$ for all $\lambda > 0$. Since p and λp lie on the same ray through the origin, distances from any fixed site decompose additively along this ray (by the same reasoning as Lemma 6.2.6). Thus:

$$F_C(a, p) = F_C(a, \lambda p) + F_C(\lambda p, p) \quad \text{and} \quad F_C(b, p) = F_C(b, \lambda p) + F_C(\lambda p, p).$$

Since $F_C(a, p) = F_C(b, p)$ by hypothesis, we have

$$F_C(a, \lambda p) + F_C(\lambda p, p) = F_C(b, \lambda p) + F_C(\lambda p, p).$$

Subtracting $F_C(\lambda p, p)$ from both sides yields $F_C(a, \lambda p) = F_C(b, \lambda p)$. Since for any point p on the bisector, $F_C(a, p) = F_C(b, p)$, we can see that the entire ray from the origin through p lies on the

bisector.

For the reverse Funk it can analogously be shown that if $F_C^r(a, p) = F_C^r(b, p)$, then $F_C^r(a, \lambda p) = F_C^r(b, \lambda p)$ for all $\lambda > 0$. $\square$

Theorem 6.2.8. *Calculating the forward and reverse Funk Voronoi diagrams of a set of sites in a cone $C \subset \mathbb{R}^d$ is equivalent to computing the forward and reverse Funk weighted Voronoi diagrams on any $(d - 1)$ -dimensional bounded cross-sectional convex region Ω .*

Proof. Since bisectors are composed of rays through the origin (see Theorem 6.2.7), the Voronoi diagram on a bounded cross section Ω of the cone is equivalent to calculating the Voronoi diagram on the entire cone. Sites are projected through the origin onto Ω and given additive weights determined by their scaling factor (see Lemma 6.2.6). A point p belongs to the Voronoi cell of a site x in C if and only if its projection onto Ω belongs to the weighted Voronoi cell of the projected site (a consequence of Lemma 6.2.6). $\square$

6.2.3.2 Bisectors in 2-Dimensional Cones

We provide explicit equations for bisectors in 2-dimensional cones. This will provide geometric intuition for the higher-dimensional constructions.

Lemma 6.2.9. *Given two sites $p, q \in \text{int } C \subset \mathbb{R}^2$ such that neither dominates the other:*

1. *The forward Funk bisector is a ray from the origin through $\partial C_p \cap \partial C_q$.*
2. *The reverse Funk bisector is a ray from the origin through $\partial C_p^r \cap \partial C_q^r$.*

Proof. Consider the forward Funk. Let the boundary rays of C be R_1 and R_2 , given by the equations $R_1 : c_1x + c_2y = 0$ and $R_2 : c_3x + c_4y = 0$. Let $p = (x_1, y_1)$ and $q = (x_2, y_2)$ lie in angular

order $\langle R_1, p, q, R_2 \rangle$. This ordering determines that distances from p are calculated using R_1 , while distances from q are calculated using R_2 .

We define the linear forms $L_1(x, y) = c_1x + c_2y$ and $L_2(x, y) = c_3x + c_4y$ corresponding to R_1 and R_2 respectively. Note that $L_i(x, y)$ gives a value proportional to the signed distance from (x, y) to ray R_i . A point (x, y) lies on the Funk bisector if and only if $F_C(p, (x, y)) = F_C(q, (x, y))$. By the similar triangles formula, this equality becomes (we can avoid absolute values because of the angular order and exponentiate to remove logarithms):

$$\frac{L_1(x_2, y_2)}{L_1(x, y)} = \frac{L_2(x_1, y_1)}{L_2(x, y)}.$$

Cross-multiplying yields:

$$L_1(x_2, y_2) \cdot L_2(x, y) = L_2(x_1, y_1) \cdot L_1(x, y).$$

Or equivalently:

$$L_2(x, y) \cdot L_1(x_2, y_2) - L_1(x, y) \cdot L_2(x_1, y_1) = 0.$$

Since this equation is linear in x and y , the bisector is a line. This line passes through both the origin (the apex of C), since the linear forms have no constant term, and the intersection point $\partial C_p \cap \partial C_q$, since it is the intersection of the zero radius Funk balls.

For reverse Funk the situation is analogous with C^r in place of C and the fractions inverted.

The resulting formula is as follows:

$$L_2(x, y) \cdot L_1(x_1, y_1) - L_1(x, y) \cdot L_2(x_2, y_2) = 0.$$

□

Note that Voronoi diagrams in both 2-dimensional Funk geometries are equivalent to angular sorting after pruning dominated sites.

6.2.4 Voronoi Diagrams in d -Dimensional Elliptical Cones

We now extend our results to d -dimensional elliptical cones. Since the *Funk metrics are affine invariants*, we will consider the circular case, to which the elliptical case can be affinely mapped [4]. We show that we can compute the forward and reverse Funk d -dimensional Voronoi diagrams in the circular cones by computing the Apollonius diagram on a set of disks representing the zero balls of the sites in $(d - 1)$ -dimensions with additional linear time filtering and relocation. *We focus on the forward Funk in this section, but the reverse Funk is analogous except when noted.*

Consider a d -dimensional circular cone $C \subset \mathbb{R}^d$ with apex at the origin and a set of sites $S = \{s_1, \dots, s_n\}$ in the interior of C . By Theorem 6.2.7 the bisectors in the Funk metrics consist of rays through the origin.

To begin, we consider the sites by Euclidean distance from the apex and let Ω be an orthogonal cut to the cone's axis below the lowest site for the reverse Funk and above the highest for the forward Funk. Note that Ω is a $(d - 1)$ -dimensional ball.

Constructing Apollonius Sites: For each site $s_i \in S$, let $B_i = C_{s_i} \cap \Omega$, the ball made from intersecting the Funk zero radius ball of s_i with Ω . For each s_i we define the corresponding Apollonius site (c_i, w_i) where c_i is the Euclidean center of ball B_i with Euclidean radius w_i . Note the Euclidean center and radius are different from the Funk center and radius. Denote this new set of sites in the Apollonius as S' .

Structural Equivalence: We now establish that the combinatorial structure of the forward

and reverse Funk Voronoi diagrams are determined by the Apollonius diagram.

Lemma 6.2.10. *If p lies on the bisector of two sites s_i and s_j in the forward Funk Voronoi diagram, then the ball around p of radius $F_C^r(p, s_i) = F_C^r(p, s_j)$ is tangent to B_i and B_j .*

Proof. This follows from the fact that $C_{F^r(p, s_i)p}$'s apex is at $\partial C_{s_i} \cap \partial C_{s_j}$ and therefore must trace the outside of the balls where they touch (see Figure 6.2 (a)). □

Corollary 6.2.11. *If a point p is on the bisector of two sites s_i and s_j in the reverse Funk Voronoi diagram then the ball around p with radius $F_C(p, s_i) = F_C(p, s_j)$ is tangent to B_i and B_j .*

Proof. This follows as Lemma 6.2.10, except the cones are reflected (see Figure 6.2 (b)). □




Figure 6.2: (a) Reverse Funk ball tangent to two forward Funk balls, (b) forward Funk ball tangent to two reverse Funk balls.

Lemma 6.2.12. *Let v' be a vertex in the Apollonius Voronoi diagram determined by sites $(c_1, w_1), \dots, (c_d, w_d)$, and let B be the corresponding ball tangent to the d Apollonius sites. Then v' corresponds to a vertex in reverse Funk if and only if $B \subset \Omega$.*

Proof. If B is not in Ω then it does not correspond to a forward Funk ball around a vertex v and hence it cannot be equidistant to the sites.

If B is in Ω then consider the cone C' congruent to C whose cross section with Ω is B . Let its apex be v . Then the zero ball of v is exactly B and since B is tangent to the zero balls of $s_1, s_2, \dots, s_d$, we know v is equidistant to those sites and v is a reverse Funk vertex. $\square$

Lemma 6.2.13. *If d sites $s_1, \dots, s_d \in S$ determine a Voronoi vertex in the forward Funk metric in Ω , then the corresponding sites in the Apollonius $(c_1, w_1), \dots, (c_d, w_d)$ determine a vertex in the Apollonius diagram.*

Proof. A point p is a vertex in the forward Funk metric between d sites $s_1, s_2, \dots, s_d$ if and only if the reverse Funk ball around p with corresponding radius is tangent to the zero radius balls around $s_1, s_2, \dots, s_d$. Call this ball B . Note that B is therefore tangent to $(c_1, w_1), \dots, (c_d, w_d)$, and hence the Euclidean center of B is equidistant in the Apollonius. $\square$




Figure 6.3: (a) The apex of C' lies in C . (b) Locating the reverse Funk vertex v in B .

Lemma 6.2.14. *Let v' be a vertex in the Apollonius Voronoi diagram determined by sites $(c_1, w_1), (c_2, w_2), \dots, (c_d, w_d)$, and let B be the corresponding ball tangent to the d Apollonius sites. Then v' corresponds to a vertex in the forward Funk Voronoi diagram.*

Proof. Let v' be an Apollonius vertex with corresponding ball B . Since B is tangent to d zero balls contained in Ω , the ball B intersects Ω . Consider the cone through B , C' , that is congruent to C .

Note that for any point p on $\partial C' \cap \Omega$ the boundary ray of C through p extending upwards must lie entirely inside C (see Figure 6.3 (a)). Since the ray emanates from the apex of C' the apex of C' must be in C . Hence C' corresponds to the zero ball around some point in C which must therefore be a Voronoi vertex in the forward Funk Voronoi. $\square$

Corollary 6.2.15. *In the reverse Funk metric, d points have a Voronoi vertex only if they have one in the forward Funk metric.*

Proof. Apollonius vertices are reverse Funk Voronoi vertices only if their forward Funk balls are subsets of Ω as opposed to just intersecting Ω (see Lemma 6.2.14). $\square$

Algorithm. Let Ω be a bounded cut above the sites in forward Funk or below for reverse Funk. We compute the forward Funk or reverse Funk Voronoi diagram as follows:

1. **Setup:** For each site $s_i \in S$, compute the zero ball $B_i = C_{s_i} \cap \Omega$. Extract the Euclidean center c_i and radius w_i of B_i to form the Apollonius site (c_i, w_i) . Let S' denote the resulting set of Apollonius sites.
2. **Apollonius:** Compute the Apollonius diagram of S' in $(d - 1)$ dimensions with complexity $O(n \log h)$ (where h is the number of nondominated sites) when $d - 1 = 2$ [47] or $O(n^{\lceil \frac{d-1}{2} \rceil + 1})$ when $d - 1 > 2$ [48]. Let V' denote the set of Apollonius vertices.
3. **Relocate:** For each Apollonius vertex $v' \in V'$, we compute the corresponding Funk vertex v by doing the following: Consider the corresponding ball B and the cone congruent to C^r (C in reverse Funk) passing through it. Calculate the cone's height and position a point at its apex. Then consider the ray through that point to the apex of C . Where that ray intersects B is where the Funk vertex must be. In the forward Funk it always exists, but in the reverse it only exists if $B \subset \Omega$ (see Lemmas 6.2.13 and 6.2.12) (see Figure 6.3 (b)).

Theorem 6.2.16. *The forward and reverse Funk Voronoi diagrams of n sites in a d -dimensional circular cone can be computed in $O(T_{d-1}(n))$ time, where $T_k(n)$ is the time to compute an Apollonius diagram of n sites in k dimensions.*

Proof. In our setup we intersect planes and cones, which is purely algebraic and $O(1)$ per site, then we prepare our data, which is constant time per site so $O(n)$. The Apollonius diagram is computed in time $O(T_{d-1}(n))$. Note there are at most $O(T_{d-1}(n))$ vertices [49]. For each of these, we require only $O(1)$ to relocate through a few simple algebraic steps. $\square$

Corollary 6.2.17. *The Voronoi diagram of a set of sites in an elliptical cone in the Funk metrics can be computed in time $O(n \log h)$ when $d - 1 = 2$, and $O(n^{\lceil \frac{d-1}{2} \rceil + 1})$ otherwise.*

Proof. This follows directly from Theorem 6.2.16 and existing Apollonius algorithms [47, 48]. $\square$

6.2.5 Voronoi Diagrams in 3-Dimensional Polygonal Cones.

For polygonal cones we first characterize the properties of the forward and reverse Funk Voronoi diagram, then we describe how to compute some primitives and show that it can be solved with the abstract Voronoi diagram framework. As we did before, we consider the sites by Euclidean distance from the apex. Let Ω be an orthogonal cut to the cone's axis below the lowest site for reverse Funk and above the highest for forward Funk. Note that Ω is a 2-dimensional polygon. Projecting our sites to Ω gives them weights. To make the Voronoi diagram we will use a modified version of the concept of spokes from the Hilbert metric [24].

Definition 6.2.18 (Forward Spokes). For a point p inside a polygon Ω with vertices $v_1, \dots, v_m$, the *forward spokes* from p are the rays from p to v_i . These spokes partition Ω into sectors around p (see Figure 6.4 (a)).

Definition 6.2.19 (Reverse Spokes). For a point p inside a polygon Ω with vertices $v_1, \dots, v_m$, the *reverse spokes* from p are the rays from p in the direction away from v_i . These spokes partition Ω into sectors around p (see Figure 6.4 (b)).




Figure 6.4: (a) Spokes in a forward Funk Voronoi and, (b) reverse Funk Voronoi.

Lemma 6.2.20. *Given two weighted sites $(s_1, w_1), (s_2, w_2)$ in a polygon Ω , the forward and reverse Funk metric bisectors between them are piecewise composed of line segments separated by sector.*

Proof. To start, we consider the forward Funk. The spokes from s_1 and s_2 partition Ω into finitely many regions. Within each region, the forward Funk distances to points from both s_1 and s_2 are computed using fixed boundary edges E_i and E_j of Ω . Let the sites and the edges lie in order $\langle E_i, s_1, s_2, E_j \rangle$. Let E_i and E_j be given by the linear forms $E_i(x, y) = 0$ and $E_j(x, y) = 0$. A point (x, y) lies on the weighted bisector if $F_C(s_1, (x, y)) + w_1 = F_C(s_2, (x, y)) + w_2$. By the similar triangles formula:

$$\ln \left(\frac{E_j(s_1)}{E_j(x, y)} \right) + w_1 = \ln \left(\frac{E_i(s_2)}{E_i(x, y)} \right) + w_2.$$

Which becomes:

$$E_j(s_1) \cdot E_i(x, y) = e^{w_2 - w_1} \cdot E_i(s_2) \cdot E_j(x, y).$$

Which is linear in (x, y) . Therefore, the bisector in each sector is a line segment (or empty if it does not intersect the sector). By the continuity of $F_C(s_1, (x, y)) + w_1 - F_C(s_2, (x, y)) - w_2$ the bisector must be one continuous curve. So we have a piecewise linear bisector overall.

In the reverse Funk the proof is analogous except we use reverse spokes and get:

$$\ln \left(\frac{E_i(x, y)}{E_i(s_1)} \right) + w_1 = \ln \left(\frac{E_j(x, y)}{E_j(s_2)} \right) + w_2.$$

Which becomes:

$$E_j(s_2) \cdot E_i(x, y) = e^{w_2 - w_1} \cdot E_i(s_1) \cdot E_j(x, y). \quad \square$$

We claim that the bisectors between two sites are composed of at most m sides. This will allow us to bound the complexity of our algorithm.

Lemma 6.2.21. *Given two weighted sites $(s_1, w_1), (s_2, w_2)$ in a polygon Ω with m vertices, the forward and reverse Funk bisectors consist of $O(m)$ line segments.*

Proof. Each of our sites has m spokes coming out of it (going towards vertices for forward Funk and away from vertices for reverse Funk). These intersect to form sectors. In each sector the bisectors are line segments that are composed piecewise to make a connected curve. However, once a bisector passes over a spoke it can never pass back over it without violating the star-shapedness of Voronoi cells.

Since each site contributes at most m spokes, the bisector crosses at most $2m$ spoke boundaries, and changes equation at most $2m$ times. So the bisector has complexity $O(m)$. $\square$

Now we can solve for the Voronoi diagram with the abstract Voronoi diagram framework [40, 50–52]. In particular we will draw from “Abstract Voronoi diagram revisited” [52]. First we will show that the forward and reverse Funk metrics satisfy the required axioms.

Axiom A1 (Jordan Curves): By Lemma 6.2.21, each bisector consists of $O(m)$ line segments forming a piecewise-linear curve in Ω . For every bisector we extend where they hit $\partial\Omega$ toward infinity. The resulting curve is piecewise linear in $\mathbb{R}^2$, unbounded, and projects to a Jordan curve through the north pole under stereographic projection.

Axiom A2 (Path-connected regions): By Lemma 6.2.5 inside Ω , each Voronoi cell is star-shaped with respect to its site and hence path-connected. Then bisectors are extended as rays outside of Ω , the portion of the cell outside Ω becomes bounded by half-planes and the boundary of Ω , and hence is path-connected. Since the interior portion is path connected, and the exterior is path-connected, and they share a boundary, their union is path connected.

Axiom A3 (Coverage): Inside Ω , every point is at finite weighted Funk distance from all sites, hence belongs to the closure of at least one Voronoi region. Outside of Ω , the extended bisector rays partition $\mathbb{R}^2 - \Omega$ into regions, each with an associated site. Therefore, the closures of all Voronoi regions cover $\mathbb{R}^2$.

To get the exact running time we must now solve the necessary primitive for our chosen abstract Voronoi diagram framework [52]. This requires us to solve for the Voronoi diagram of five sites as a primitive. Then “Abstract Voronoi diagram revisited” [52] gives our runtime.

Lemma 6.2.22. *The Funk Voronoi diagrams of five sites can be computed in time $O(m)$.*

Proof. We apply a modified version of the randomized incremental algorithm from previous work on Voronoi diagrams in the Hilbert metric [31] but on five sites. We calculate the Funk distance

between all pairs of sites and remove dominated ones. The algorithm inserts sites one at a time, tracing the boundary of each new Voronoi cell with spokes (we only need forward spokes or reverse spokes depending on the Funk metric) through bisector segments and sector boundaries. For a constant number of sites (five in this case), the diagram has $O(m)$ complexity (refer to Section 6.2.6 to see that three sites contribute at most two Voronoi vertices) with $O(m)$ total bisector segments and construction time $O(m)$. $\square$

Theorem 6.2.23. *The forward or reverse Funk Voronoi diagram of n sites in a 3-dimensional polygonal cone with m facets can be computed in $O(nm \log(n))$ time.*

6.2.6 Circumcenter Characterization in 3-Dimensional Cones

One particular aspect of the Funk metrics is that three points in the Funk metrics do not necessarily lie on the boundary of a ball. The characterization with respect to both Funk metrics is almost entirely the same. We will make our definitions direction agnostic. Given three sites we characterize when they admit a forward or reverse Funk circumcenter in 3-dimensional cones. This analysis applies to any convex cone.

Let p, q, r be three sites in C . Without loss of generality assume that r is the site furthest from the apex of the cone (or closest in the reverse Funk). We assume that neither p dominates q nor q dominates p . Let H be a plane through r that is perpendicular to the axis of the cone. Let $\Omega = H \cap C$. By Theorem 6.2.8 we obtain weighted sites p' with weight w_p and q' with weight w_q on Ω . Since r lies on the plane, it requires no weight.

Lemma 6.2.24. *If p, q, r admit a circumcenter u in C in the forward or reverse Funk metrics, then, every site along the ray through the apex and u is the center of a ball through p, q, r .*

Proof. This follows from Theorem 6.2.7. □

Hence we can determine circumcenter existence in planar cross sections of C with weighted point-sites rather than the full cone. We introduce our first region.

Definition 6.2.25 (Inner Regions: Z_0). Define the zero balls $Z_p = C_p \cap \Omega$ and $Z_q = C_q \cap \Omega$, and let $Z_0 = Z_p \cup Z_q$.

These are the weighted zero radius balls of p and q in the plane.

Lemma 6.2.26. *If $r \in \text{int } Z_0$ then r has no cell in the Voronoi diagram.*

Proof. If r is in Z_p then r is dominated by p (see Lemma 6.2.3 and Lemma 6.2.1). □

For the remainder of the section assume $r \notin Z_0$. The weighted $((p', w_p), (q', w_q))$ -bisector intersects $\partial\Omega$ at two points u and v in either Funk metric (by Lemma 6.2.5). We now define the tangent balls at these points (equivalent to infinite balls in Hilbert [33]).

Definition 6.2.27 (Tangent Balls). Let $u, v \in \partial\Omega$ be the two points where the $((p', w_p), (q', w_q))$ -bisector intersects $\partial\Omega$ such that p', u, q' is in clockwise order around Ω and p', v, q' is in counter-clockwise order. Define $B(p : q)$ to be the Funk ball centered at u tangent to both Z_p and Z_q , and $B(q : p)$ to be the Funk ball centered at v tangent to both.

Both these Funk balls must be in the direction opposite of the bisectors incident direction. We define the overlap region to be the intersection of the two tangent balls (see Figure 6.5 (a)).

Definition 6.2.28 (Overlap Region). Given two sites $p, q \in \text{int } C$, the overlap region in the Funk metrics is the intersection of both tangent balls at $\partial\Omega$: $Z_1(p, q) = B(p : q) \cap B(q : p)$.

Lemma 6.2.29. *If r is in $Z_1(p, q)$ then p, q, r have no circumcenter.*

Proof. The $((p', w_p), (q', w_q))$ -bisector parameterizes centers of balls tangent to Z_p and Z_q within Ω . The balls $B(p : q)$ and $B(q : p)$ are the limiting tangent balls as the center approaches $\partial\Omega$. Any ball passing through a point in $Z_1 = B(p : q) \cap B(q : p)$ would have to have its center outside of Ω . Hence r lies interior to all tangent balls to Z_p and Z_q . $\square$

Next we define the one center regions, then two center regions.

Definition 6.2.30 (One-Center Region). Define $B_1(p, q) = (B(p : q) \cup B(q : p)) - Z_1$.

Lemma 6.2.31. *If r lies within $B_1(p, q)$ then there is one ball through p, q and r .*

Proof. Without loss of generality let $r \in B(p : q)$ and $r \notin B(q : p)$. Then the pencil of Funk balls from v to u passes from one side of r to the other. By the intermediate value theorem, there is a Funk ball passing through p, q, r in Ω . By continuity, it passes through once. $\square$

Definition 6.2.32 (Two Center Region). Given two sites $p, q \in \text{int } C$, the two center region is the space between $Z_p, Z_q, B(p : q), B(q : p)$. We call it $Z_2(p, q)$.

Lemma 6.2.33. *If $r \in Z_2(p, q)$ then it lies on the boundary of two Funk balls through p and q .*

Proof. This region is the region traced out by the boundaries of the Funk balls (in the opposite direction to the Voronoi) outside of $B(p : q)$ and $B(q : p)$. This region can come in three forms: there is a Z_1 region and this Z_2 region is split into two pieces on either side of it, there is no Z_1 region and the Z_2 region is one connected piece, there is no Z_1 region and Z_p and Z_q intersect and the Z_2 is split above and below them (see Figure 6.5 (b)-(e)). We will tackle all these cases at once. Consider that by the nature of r not being in either infinite ball then p and q are both closer to the boundary than r so r 's cell is sandwiched between them. This means there are at least two

circumcenters. However by star-shapedness of the cells we can see that we can leave and enter the cell at most once each and so there are exactly two circumcenters. $\square$




Figure 6.5: In the reverse Funk, (a) Cross section of a circular cone, (b) (c) (d) regions labeled in a circular cone with differently weighted sites, (e) regions labeled in a polygonal cone, (f) pencil of balls along the bisector in a polygonal cone.

Now we define our final region, the outer region. This is the area exterior to all the other regions, on the outside of $B(p : q)$, $B(q : p)$, Z_p , and Z_q (see Figures 6.5(b)-(e)).

Definition 6.2.34 (Outer Region). The outer region of p and q , $W(p, q)$ is $\Omega - (Z_0(p, q) \cup Z_1(p, q) \cup Z_2(p, q) \cup B(p : q) \cup B(q : p))$

Lemma 6.2.35. *If $r \in W(p, q)$ then there is no ball through r , p , and q .*

Proof. In this case r is outside of the two infinite balls $B(p : q)$ and $B(q : p)$ and $Z_2(p, q)$. Hence no ball on the pencil of balls on the bisector will pass through it (see Figure 6.5 (f)). $\square$

Theorem 6.2.36. *Let p, q be two sites in a 3-dimensional cone C with $r \in C$ such that neither p nor q dominates each other in the forward Funk metric (or reverse analogously). Then the number of forward Funk circumcenters (or reverse) through p, q, r is determined by which region r occupies:*

- *If $r \in Z_0(p, q)$, there are 0 circumcenters (Lemma 6.2.26).*
- *If $r \in Z_1(p, q)$, there are 0 circumcenters (Lemma 6.2.29).*
- *If $r \in B_1(p, q)$, there is exactly 1 circumcenter (Lemma 6.2.31).*
- *If $r \in Z_2(p, q)$, there are exactly 2 circumcenters (Lemma 6.2.33).*
- *If $r \in W(p, q)$, there are 0 circumcenters (Lemma 6.2.35).*

Chapter 7: Conclusions

In this dissertation, we have provided the foundations for further study of the Hilbert and Funk metrics in computational geometry. We began with an introduction in Chapter 1 outlining our motivation for studying these metrics, then in Chapter 2 we provided an overview of the basic tools needed to study the Hilbert and Funk metrics.

In Chapter 3, we presented the first algorithmic study of the Hilbert metric in a convex m sided polygon. We provided two randomized incremental algorithms and one deterministic divide and conquer for Voronoi diagrams in the Hilbert metric. The first randomized algorithm and the deterministic algorithm computed the Voronoi diagram of a set of n point sites. The second randomized algorithm extended this to compute the Voronoi diagram of the set of n sites, each of which was allowed to be a point site or a line segment site. Our algorithms ran in expected time $O(mn \log n)$. The algorithms used $O(mn)$ storage, which matched the worst-case combinatorial complexity of the Voronoi diagram in the Hilbert metric.

In Chapter 4, we studied the properties of Hilbert bisectors and analyzed Hilbert Voronoi diagrams on point sites in the dynamic setting. We provided explicit equations for Hilbert bisectors and found that they were composed of piecewise conics. We then identified the forms of the conics that the bisectors could manifest, and identified when they degenerated into straight lines. We also identified the discontinuities that caused bisectors to become two dimensional.

In Chapter 5, we showed how to adapt the Euclidean Delaunay triangulation to the Hilbert geometry on convex polygons in the plane. We presented a simple randomized incremental algorithm that computed the Hilbert Delaunay triangulation (which we found was not the same as the Euclidean) for a set of n points in the Hilbert geometry defined by a convex m -gon. The algorithm ran in $O(n(\log n + \log^3 m))$ expected time. In addition, we introduced the notion of the Hilbert hull of a set of points, which we defined to be the region covered by their Hilbert Delaunay triangulation. We presented an algorithm for computing the Hilbert hull in time $O(nh \log^2 m)$, where h is the number of points on the hull's boundary.

In Chapter 6, we presented the first algorithmic study of the Funk metrics in the conical geometry. Our main contributions included: establishing that bisectors consisted of rays through the apex of the cone; proving a dimensional reduction from d -dimensional Funk Voronoi diagrams to $(d - 1)$ -dimensional additively weighted Voronoi diagrams on bounded cross sections of the cone; providing efficient algorithms for computing these in d -dimensional elliptical cones in time $O(n^{\lceil \frac{d-1}{2} \rceil + 1})$ and 3-dimensional polygonal cones with m sides in time $O(mn \log(n))$; and categorizing circumcenters for 3-dimensional cones.

While our results provide a strong foundation for studying the Hilbert and Funk metrics, there is still a considerable amount of future work to be done in investigating the geometry of these metrics and general metrics defined within convex bodies. We list some open problems here according to perceived accessibility, from most challenging to most accessible.

- Extending algorithms in the Hilbert metric to d -dimensional convex polytopes.
- Computing minimum perimeter enclosing balls of a set of sites in the Hilbert metric.
- Computing weighted Voronoi diagrams in the Hilbert metric.

- Investigating computational geometry in the Thompson metric (the maximum of the forward and reverse Funk metrics).
- Extending algorithms in the Hilbert metric to convex bodies which have both smooth and flat edges such as the correlation ellipsope.
- Delaunay triangulations in the Funk metrics on d -dimensional elliptical cones and 3-dimensional polygonal cones.
- Computing furthest point Voronoi diagrams in the d -dimensional Funk metrics.

Bibliography

- [1] D. Hilbert. Ueber die gerade Linie als kürzeste Verbindung zweier Punkte. *Math. Annalen*, 46:91–96, 1895.
- [2] Frank Nielsen and Ke Sun. Clustering in Hilbert’s projective geometry: The case studies of the probability simplex and the ellipsope of correlation matrices. In Frank Nielsen, editor, *Geometric Structures of Information*, pages 297–331. Springer Internat. Pub., 2019.
- [3] Frank Nielsen and Ke Sun. Non-linear embeddings in Hilbert simplex geometry, 2022.
- [4] Athanase Papadopoulos and Marc Troyanov. *Handbook of Hilbert geometry*, volume 22 of *IRMA Lectures in Mathematics and Theoretical Physics*. European Mathematical Society Publishing House, 2014.
- [5] Sunil Arya, Guilherme Dias da Fonseca, and David M. Mount. Approximate polytope membership queries. *SIAM J. Comput.*, 47(1):1–51, 2018.
- [6] Ahmed Abdelkader, Sunil Arya, Guilherme Dias da Fonseca, and David M. Mount. Approximate nearest neighbor searching with non-Euclidean and weighted distances. In *Proc. 30th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 355–372, 2019.
- [7] Sunil Arya, Guilherme Dias da Fonseca, and David M. Mount. Near-optimal ε -kernel construction and related problems. In *Proc. 33rd Internat. Sympos. Comput. Geom.*, pages 10:1–15, 2017.
- [8] Friedrich Eisenbrand, Nicolai Hähnle, and Martin Niemeier. Covering cubes and the closest vector problem. In *Proc. 27th Annu. Sympos. Comput. Geom.*, pages 417–423, 2011.
- [9] Thomas Rothvoss and Moritz Venzin. Approximate CVP in time $2^{0.802n}$ – Now in any norm! In *Proc. 23rd Internat. Conf. on Integ. Prog. and Comb. Opt. (IPCO 2022)*, pages 440–453, 2022.
- [10] Friedrich Eisenbrand and Moritz Venzin. Approximate CVPs in time $2^{0.802n}$. *J. Comput. Sys. Sci.*, 124:129–139, 2021.
- [11] Márton Naszódi and Moritz Venzin. Covering convex bodies and the closest vector problem. *Discrete Comput. Geom.*, 67:1191–1210, 2022.
- [12] Sunil Arya, Guilherme Dias da Fonseca, and David M. Mount. On the combinatorial complexity of approximating polytopes. *Discrete Comput. Geom.*, 58(4):849–870, 2017.

- [13] Rahul Arya, Sunil Arya, Guilherme Dias da Fonseca, and David M. Mount. Optimal bound on the combinatorial complexity of approximating polytopes. *ACM Trans. Algorithms*, 18:1–29, 2022.
- [14] Constantin Vernicos. On the Hilbert geometry of convex polytopes. In *Handbook of Hilbert geometry*, volume 22 of *IRMA Lectures in Mathematics and Theoretical Physics*, pages 111–126. European Mathematical Society Publishing House, 2014.
- [15] Ahmed Abdelkader and David M. Mount. Economical Delone sets for approximating convex bodies. In *Proc. 16th Scand. Workshop Algorithm Theory*, pages 4:1–4:12, 2018.
- [16] Ahmed Abdelkader and David M Mount. Convex approximation and the hilbert geometry. In *2024 Symposium on Simplicity in Algorithms (SOSA)*, pages 286–298. SIAM, 2024.
- [17] Constantin Vernicos and Cormac Walsh. Flag-approximability of convex bodies and volume growth of Hilbert geometries, 2018.
- [18] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná. Hyperbolic geometry of complex networks. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 82(3):036106, 2010.
- [19] Ginestra Bianconi and Christoph Rahmede. Emergent hyperbolic network geometry. *Scientific Reports*, 7(1):41974, 2017.
- [20] Weihua Yang and David Rideout. High dimensional hyperbolic geometry of complex networks. *Mathematics*, 8(11):1861, 2020.
- [21] David Reeb, Michael J. Kastoryano, and Michael M. Wolf. Hilbert’s projective metric in quantum information theory. *J. Math. Physics*, 52(8), 2011.
- [22] Bas Lemmens and Roger Nussbaum. Birkhoff’s version of Hilbert’s metric and its applications in analysis, 2013.
- [23] Yongxin Chen, Tryphon Georgiou, and Michele Pavon. Entropic and displacement interpolation: A computational approach using the Hilbert metric. *SIAM J. Appl. Math.*, 76:2375–2396, 2016.
- [24] Frank Nielsen and Laetitia Shao. On balls in a Hilbert polygonal geometry. In *Proc. 33rd Internat. Sympos. Comput. Geom.*, pages 67:1–67:4, 2017. (Multimedia contribution).
- [25] Paul Funk. Über Geometrien, bei denen die geraden die kürzesten sind. *Mathematische Annalen*, 101(1):226–237, 1929.
- [26] Anthony C. Thompson. On certain contraction mappings in a partially ordered vector space. *Proceedings of the American Mathematical Society*, 14(3):438–443, 1963.
- [27] Dmitry Faifman. A funk perspective on billiards, projective geometry and mahler volume. *Journal of Differential Geometry*, 127(1):161–212, 2024.

- [28] Athanase Papadopoulos and Marc Troyanov. Weak finsler strutures and the funk metric. *arXiv preprint arXiv:0804.0705*, 2008.
- [29] Zhongmin Shen. Funk metrics and r-flat sprays. *arXiv preprint math/0109037*, 2001.
- [30] Newton Mayer Solórzano Chávez, Víctor Arturo Martínez León, Alexandre Henrique Rodrigues Filho, and Marcelo Almeida de Souza. Geometry of a navigation problem: The λ -Funk Finsler metrics. *Ciência e Natura*, 47:e88467–e88467, 2025.
- [31] Auguste H. Gezalyan and David M. Mount. Voronoi diagrams in the Hilbert metric. In *39th International Symposium on Computational Geometry (SoCG 2023)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2023.
- [32] Madeline Bumpus, Caesar Dai, Auguste H. Gezalyan, Sam Munoz, Renita Santhoshkumar, Songyu Ye, and David M. Mount. Analysis of dynamic Voronoi diagrams in the Hilbert metric. In *Proc. 35th Canadian Conf. Comput. Geom. (CCCG 2023)*, Montreal, Canada, 2023.
- [33] Auguste H Gezalyan, Soo H Kim, Carlos Lopez, Daniel Skora, Zofia Stefankovic, and David M Mount. Delaunay triangulations in the hilbert metric. In *19th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2024)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.
- [34] Herbert Busemann. *The Geometry of Geodesics*. Academic Press, 1955.
- [35] Athanase Papadopoulos and Marc Troyanov. From Funk to Hilbert geometry. In *Handbook of Hilbert geometry*, volume 22 of *IRMA Lectures in Mathematics and Theoretical Physics*, pages 33–68. European Mathematical Society Publishing House, 2014.
- [36] Franz Aurenhammer. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM Computing Surveys (CSUR)*, 23(3):345–405, 1991.
- [37] Steven Fortune. Voronoi diagrams and Delaunay triangulations. In *Handbook of discrete and computational geometry*, pages 705–721. Chapman and Hall/CRC, 2017.
- [38] Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 3rd edition, 2010.
- [39] Marc Troyanov. Funk and Hilbert geometries from the Finslerian viewpoint. In *Handbook of Hilbert geometry*, volume 22 of *IRMA Lectures in Mathematics and Theoretical Physics*, pages 69–110. European Mathematical Society Publishing House, 2014.
- [40] K. Mehlhorn, St. Meiser, and C. Ó’Dúnlaing. On the construction of abstract Voronoi diagrams. *Discrete Comput. Geom.*, 6:211–224, 1991.
- [41] Kenneth L. Clarkson and Peter W. Shor. Applications of random sampling in computational geometry, ii. *Discrete Comput. Geom.*, 4:387–421, 1989.
- [42] Franco P. Preparata and Michael Ian Shamos. *Computational Geometry: An Introduction*. Springer, 1985.

- [43] Godfried T. Toussaint. The relative neighbourhood graph of a finite planar set. *Pattern Recogn.*, 12:261–268, 1980.
- [44] Leonidas J. Guibas, Donald E. Knuth, and Micha Sharir. Randomized incremental construction of Delaunay and Voronoi diagrams. *Algorithmica*, 7:381–413, 1992.
- [45] Ray A Jarvis. On the identification of the convex hull of a finite set of points in the plane. *Information Processing Letters*, 2(1):18–21, 1973.
- [46] Sunil Arya and David M Mount. Cauchy’s surface area formula in the funk geometry. *arXiv preprint arXiv:2601.16468*, 2026.
- [47] Menelaos I Karavelas and Mariette Yvinec. Dynamic additively weighted Voronoi diagrams in 2d. In *European Symposium on Algorithms*, pages 586–598. Springer, 2002.
- [48] Jean-Daniel Boissonnat and Christophe Delage. Convex hull and Voronoi diagram of additively weighted points. In *European Symposium on Algorithms*, pages 367–378. Springer, 2005.
- [49] Jean-Daniel Boissonnat and Menelaos I Karavelas. On the combinatorial complexity of Euclidean Voronoi cells and. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete algorithms*, page 305. SIAM, 2003.
- [50] Rolf Klein, Kurt Mehlhorn, and Stefan Meiser. Randomized incremental construction of abstract Voronoi diagrams. *Comput. Geom. Theory Appl.*, 3(3):157–184, 1993.
- [51] Rolf Klein. Abstract Voronoi diagrams and their applications. In *Workshop on Computational Geometry*, pages 148–157. Springer, 1988.
- [52] Rolf Klein, Elmar Langetepe, and Zahra Nilforoushan. Abstract Voronoi diagrams revisited. *Computational Geometry*, 42(9):885–902, 2009.