The Hilbert Geometry: Algorithms and Applications

Nithin Parepally[1]
Department of Computer Science, University of Maryland, College Park, USA

Auguste H. Gezalyan[2]
Department of Computer Science, University of Maryland, College Park, USA

Abstract

The Hilbert metric is a distance function defined for points lying within a convex body. It generalizes the Cayley-Klein model of hyperbolic geometry to any convex set. This survey explores algorithms and applications of this metric, particularly in the areas of convex approximation (Macbeath Regions, flag-approximability), clustering in probability simplices, Voronoi diagrams and Delaunay Triangulations, and quantum information theory. While existing surveys focus on applications of the Hilbert metric in the context of real analysis, this work highlights its relevance in computational settings.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Approximate Membership Queries, Clustering, Voronoi diagrams, Delaunay Triangulations, Hilbert metric, Convexity Theory, Randomized Algorithms, Quantum Information Theory, Contraction Bounds

Contents

1 Introduction **1**

2 Preliminaries **2**

3 Convex approximation **4**
3.1 Economical Delone Sets. 4
3.2 Flag Approximability 5

4 Clustering **7**

5 Algorithms and Theory **10**
5.1 Voronoi Diagrams 11
5.2 Delaunay Triangulations 13

6 Quantum Information Theory **15**

7 Conclusion **18**

8 Software & Code **19**

1 Introduction

The Hilbert metric was introduced by David Hilbert in 1895 [22] as a generalization of the Cayley-Klein model of hyperbolic geometry to arbitrary convex bodies. Among its more notable properties are that straight lines are geodesics and that the Hilbert metric is preserved under projective transformations. As a result, the Hilbert metric has been applied to a variety of subjects in computer science and adjacent fields.

The Hilbert metric is an object of study in the field of convex approximation. The efficient approximation of convex bodies is a fundamental pursuit of computational geometry due to its wide range of applications including: approximate polytope membership queries [7], approximate nearest neighbor searching [1], computing approximating polytopes with low combinatorial complexity [5, 6], collision detection [38] and others.

Many techniques utilized in the approximation of high-dimensional convex bodies use an object known as a Macbeath region [2, 5, 6]. It has since been shown that these regions are equivalent to Hilbert balls up to constant factors [2]. Hilbert balls have been used directly in research involving convex approximation [3]. Furthermore, the Hilbert metric has been used in the study of the structure of convex bodies [35, 36], especially in the context of flag approximability of convex polytopes [17, 37].

There has been growing interest in applying the Hilbert metric to machine learning. Recently, the Hilbert metric has been studied in the context of SVMs [34]. Work by Frank Nielsen and Ke Sun experimented using the Hilbert metric in various contexts such as clustering points in the probability simplex [26] and graph embeddings [27] and have found it to be competitive. In particular, Nielsen's paper on clustering uses the fact that the Hilbert metric is an effective metric on discrete probability distributions in the simplex.

In quantum information theory, applications of the Hilbert metric have been realized on convex cones representing classes of matrices such as positive, separable, or PPT operators [31], where the metric is used to study contraction bounds on or between these cones. For more information on the Hilbert metric and contraction ratios, there is a prominent survey on the applications of the Hilbert metric in real analysis written by Lemmens and Nussbaum [23]. In this work however, we focus on the applications of the Hilbert metric in computer science in which there have been numerous theory works including computing Delaunay triangulations [19], Voronoi diagrams [13, 20], and SVMs [34]. Accompanying these works are several pieces of software made to study the Hilbert metric, two that allow dynamic movement of Hilbert balls and other structures [10, 24] and one that allows for the creation of Hilbert balls in Ipe extensible drawing editor [30]. We will start with some preliminaries on the Hilbert metric and then move on to some results in different fields.

2 Preliminaries

Given a fully dimensional convex body $\Omega \subset \mathbb{R}^d$, let $\operatorname{int}\Omega$ and $\partial\Omega$ denote the interior and boundary of Ω respectively. Given two points in the interior of Ω, p, q, let $\chi(p, q)$ refer to the chord through p and q in Ω. The Hilbert metric is a distance function on $\operatorname{int}\Omega$.

▶ **Definition 1** (Hilbert Metric, polytopal). *Given a fully dimensional, bounded, convex body $\Omega \subset \mathbb{R}^d$ containing two points distinct $p, q \in \operatorname{int}\Omega$, let p' and q' endpoints of $\chi(p, q)$ such that the points lie in order $\langle p', p, q, q' \rangle$ on the chord. The Hilbert metric on $\operatorname{int}\Omega$ is defined as:*

$$H_\Omega(p, q) = \frac{1}{2} \log \left(\frac{|p'q|}{|p'p|} \frac{|q'p|}{|q'q|} \right)$$

where $H_\Omega(p, p) = 0$.

The Hilbert metric on polytopes has several desirable properties: all straight lines are geodesics, though not all geodesics are straight lines; the metric is a projective invariant (it preserves the cross-ratio, a property referred to as invariance of the cross ratio); and the metric generalizes the Cayley-Klein metric, which is defined on Euclidean balls, to arbitrary convex polygons or polytopes.

To better understand the Hilbert metric, it serves to study the behavior of balls in the metric. A ball around a point, p, of a radius r in a metric space is the set of points in that space that are r away from p. In the Hilbert metric, the shape of balls around points depends on the dimensionality and complexity of Ω. Given a two-dimensional Ω with m sides, balls in the Hilbert metric are polygons with between m and $2m$ sides [24].

In two dimensions, Hilbert balls are created by extending spokes through the vertices of Ω and the center of the ball. Each wedge formed by these spokes defines a new part of the ball's boundary.

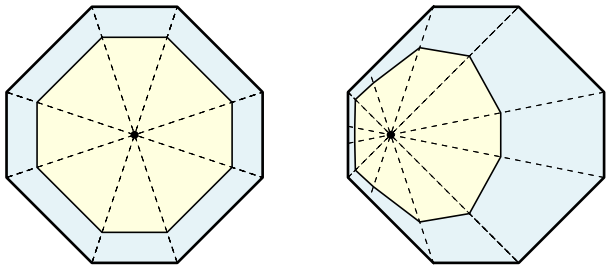

(a) (b)

Figure 1 (a) A Hilbert ball with m sides (b) A Hilbert ball with more than m sides.

Note that the Hilbert metric can be extended to points in the interior of a convex cone C, though it loses the identity of indiscernibles. That is, two points on the same ray from the vertex of the cone have zero Hilbert distance from each other. If the vertex of the cone is at the origin, then for any $\lambda \geq 0$, we have:

$$H_C(p, \lambda p) = 0$$

Alternatively, this can be viewed as a metric on rays or directional vectors originating from the vertex of the cone. In all cases in the cone, Hilbert balls remain the same: they are sub-cones C' contained within the cone C. This follows directly from the invariance of the cross-ratio under the Hilbert metric.

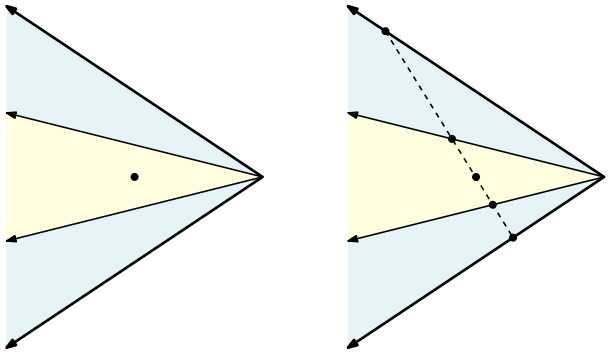

(a) (b)

Figure 2 (a) A ball in the Hilbert conical metric (b) No matter where the two points on the boundary of the ball are positioned, the Hilbert distance remains the same by projective invariance.

One of the main cases of study is the Hilbert metric in the probability simplex. A categorical distribution is a discrete probability distribution $p_1, \ldots, p_d$, on a set of events

$E_1, \ldots, E_d$. Note that $\sum_{i=1}^{d} p_i = 1$ and $p_i \geq 0$ for all i. The set of probability distributions over a set of d events is a $d-$simplex, a d-dimensional triangle. The Hilbert metric is a natural metric between these categorical distributions and is the subject of study in [26, 27].

3 Convex approximation

3.1 Economical Delone Sets.

Epsilon Approximate Membership Queries. A well known problem in computational geometry and convex approximation is the membership query problem where the goal is to determine whether a given point or item belongs to or is a "member" of a set or data structure. Exact membership queries yield a conclusion with absolute certainty, where a query result reflects the presence of a point in the set or data structure. Approximate membership queries differ in that they provide a result that is within a certain threshold of error. When there is room for error, computational complexity can be greatly reduced whilst achieving consistent results. In epsilon approximate membership queries, the parameter ϵ sets a maximum false positive error level or uncertainty that will be allowed in the query result. Applications of epsilon approximate membership queries can be seen when employing probabilistic data structures like bloom filters or cuckoo filters, quickening dictionaries [12], detecting spam, and using hash compaction functions [33].

Epsilon approximate polytope membership queries are a type of epsilon approximate membership queries. The difference arises from the fact that epsilon approximate polytope membership queries look at the presence of a query point in a bounded convex polytope. Let K be a bounded convex polytope in R^d and $diam(K)$ denote the diameter of K. Consider a query point $q \in R^d$. An epsilon approximate polytope membership query returns a positive result if $q \in k$ and a negative result if the distance between q and its closest point in $K > \epsilon \cdot diam(k)$. Note, $\epsilon \cdot diam(k)$ refers to the maximum distance allowed between q and its closest point in K to be considered within in the polytope where the real parameter ϵ represents the maximum amount of uncertainty allowed when approximating the presence of a query point in a bounded convex polytope.

Macbeath Regions & Ellipsoids. Macbeath regions are regions resulting from the intersection of a convex body K and its reflection around a point $x \in K$ referred to as the center of the region. A λ-scaled Macbeath region at a point x is defined as $M_K^{\lambda}(x) := x + \lambda((K-x) \cap (x-K))$ where λ is the factor by which the region is scaled by. These regions have an important property called *expansion-containment* where if two shrunken Macbeath regions overlap, one of the regions can be expanded appropriately such that it contains the other region. *Macbeath Ellipsoids* are maximum volume ellipsoids that are contained within a λ-scaled Macbeath region. Two interesting properties of these ellipsoids, shown by Chazelle and Matousek, are that they are unique and can be computed for any convex polytope in linear time where the time grows respective of the number of the polytope's bounding halfspaces [2].

Macbeath regions are a classical structure in convexity theory which have been used in convex set theory, geometry of numbers, the construction of lower bounds for range searching in the field of computational geometry. A particular application of these regions were explored in 2017 by Guilherme D. da Fonseca, David Mount, and Sunil Arya when they used Macbeath regions to bound the epsilon approximation of an arbitrary convex body K. This was possible through the construction of a data structure made up of nested rings within a convex body surrounding its origin. Each ring is composed of scaled Macbeath regions where if a ray is shot from the origin, at least one Macbeath region of each ring is

struck. Because ellipsoids simplify query processing times due to their nature of reasonably approximating complex shapes, the Macbeath regions were replaced with Macbeath Ellipsoids to reduce computational complexity. With each successive ring, a better approximation of the convex body boundary is reached. The farthest ring forms an epsilon approximation to the boundary [8].

Previous solutions to the approximate polytope membership problem resulted in either a logarithmic query time and a storage complexity of $O(1/\epsilon^d - 1)$ or a $O(1/e^{d-1/8})$ query time and a storage complexity of $O(1/e^{(d-1)/2})$. Mount, Arya, and Fonsecaand were able to achieve a logarithmic query time and a storage complexity of $O(1/e^{(d-1)/2})$ with their data structure. A 2018 paper by Mount and Adelkader tackles the complexity associated with creating the data structure and performing corresponding analyses derived from the ray shooting method, by offering an approach that relates Macbeath regions to Delone sets through the use of the Hilbert Metric. This approach achieves the same space-time optimal bounds discovered in 2017 through "simpler and more intuitive" means [8].

Hilbert Geometry and Approximate Polytope Membership. To construct their data structure in their 2018 paper, Mount and Adelkader viewed Macbeath regions as Delone Sets in order to leverage the packing and covering properties of Delone sets. Utilizing the Hilbert metric and properties derived from Hilbert geometry which is made up of a convex domain K in $\mathbb{R}^d$ with a Hilbert distance f_K, it was shown that shrunken Macbeath regions are nested between Hilbert metric balls. Because of this, Macbeath regions and their respective Macbeath ellipsoids act as metric balls in the Hilbert metric space up to certain constant factors of scaling.

To perform an Approximate polytope membership with Mount and Adelkader's approach, the first step is to construct a rooted layered DAG structure. This structure is built on the basis of hierarchical Delone sets, which consider sets $\{X_0 ... X_l\}$, where X_i denotes a Delone set for K_{δ_i} that is in one-to-one correspondence with the nodes at level i. These sets contain a maximal number of points within the convex body such that the packing ellipsoids are pairwise disjoint.

The nodes of DAG have two features. A cell, which is the covering ellipsoid, and a set of children (nodes belonging to the Delone set $X_i - 1$ such that the cells of the node at level i, the current node, and the node at level $i - 1$ do not overlap). The root and leaf nodes reside at levels l and 0 respectively. To determine whether a query point is a member of a bounded convex polytope with the DAG structure, starting at the root, we visit the node at each level iteratively where the point belongs to the cell of that node. Considering the case where the point does belong to the convex body, through the covering properties of Delone sets, it is known that at least one of the children of this node must contain the point. We continue to search the children of nodes for which the point lies in and end up at the leaf node whose cell contains the query point. The cell of the leaf node serves as a witness to the point's membership. In the case the point does not belong to the bounded convex polytope, the query algorithm terminates and returns a result indicating the point is not a member.

3.2 Flag Approximability

Another well known problem in convex approximation is approximating convex bodies with simple polytopes, typically considering measures such as the number of vertices, facets, and faces [37]. An alternative measure explored by Vernicos and Walsh is the number of *maximal flags* which are a finite sequence of a nested polytope faces. In particular, they introduce *flag approximaibility* as a property of a convex body describing its complexity when using

the number of maximal flags. Vernicos and Walsh show that this property directly relates to the *volume entropy* of the Hilbert geometry of the convex body. This section of the survey is based primarily on the work on Vernicos and Walsh from their 2021 paper *"Flag Approximability of Convex Bodies and Volume Growth of Hilbert Geometries"*.

Maximal Flags and Flag Approximability The *maximal flag* of a d-dimensional, closed, and convex polytope $\mathcal{P}$ is is tuple of faces $(f_0, \dots f_d)$ such that face f_i lies in dimension i and is contained within the boundary of f_{i+1}. The *flag-approximability* of a convex body Ω, $a_f(\Omega)$ is defined with respect to $N_f(\epsilon, \Omega)$ the least number of maximal flags approximating Ω within a certain Hausdorff distance,

$$a_{\mathrm{f}}(\Omega) = \liminf_{\varepsilon \to 0} \frac{\log N_{\mathrm{f}}(\varepsilon, \Omega)}{-\log \varepsilon}.$$

Volume Entropy in the Hilbert Geometry Volume entropy in the Hilbert geometry measures the rate at which the volume of a Hilbert ball grows as its radius increases. Generally, there exists upper and lower volume entropies for the Hilbert geometry of Ω which can be defined by using either the Busemann or Holmes Thompson volumes. The *lower volume entropy*, denotes the minimum growth rate of the volume of a Hilbert ball as its radius grows towards infinity or $\partial\Omega$. This quantity is described by the Holmes-Thompson Volume (or Busemann Volume) Vol^H given a Hilbert ball centered at point $p \in \mathrm{int}(\Omega)$ with radius $R > 0$ and is not affected by the position of the center.

$$\mathrm{Ent}_{lower}(\Omega) = \liminf_{R \to \infty} \frac{\log \mathrm{Vol}^{\mathrm{H}}\left(B_\Omega(p, R)\right)}{R}.$$

The *upper volume entropy* denotes the maximum growth rate of the volume of a Hilbert ball grows towards $\partial\Omega$ and replaces the infimum in the previous definition with a supremum.

$$\mathrm{Ent}_{upper}(\Omega) = \limsup_{R \to \infty} \frac{\log \mathrm{Vol}^{\mathrm{H}}\left(B_\Omega(p, R)\right)}{R}.$$

With these definitions, Vernicos and Walsh show that the lower (or upper) volume entropy of the Hilbert geometry of Ω is exactly equal to twice its flag-approximability. That is,

$$\mathrm{Ent}_{lower}(\Omega) = 2 \cdot a_{\mathrm{f}}(\Omega)$$

The proof for this involves showing the upper and lower bounds on the volume entropy with respect to the flag approximability. The upper bound by proven by first bounding the volume entropy of asymptotic balls of arbitrary radii in terms of the number of flags in a given polytope. In particular, Vernicos and Walsh show that there exists a polynomial p_d of degree d depending solely on the dimension such that the Holmes-Thompson volume of an asymptotic ball centered at an apex o with radius R is bounded by the least number of flags approximating Ω within a Hausdorff distance ϵ,

$$\mathrm{Vol}_P^H(\mathrm{As}B_P(o, R)) \leq N_f(\epsilon, \Omega) p_d(R)$$

where an asymptotic ball is a dilation of Ω about a point $y \in \Omega$ by a factor $1 - e^{2R}$.

$$\mathrm{As}B_\Omega(y, R) = y + (1 - e^{2R})(\Omega - y)$$

Then, using the relationship $B_\Omega(o, R') \subset \mathrm{As}B_\Omega(o, R') \subset \mathrm{As}B_P(o, R)$ where R' is some radius, Vernicos and Walsh conclude,

$$\frac{1}{R'} \log \mathrm{Vol}_\Omega^{\mathrm{H}}\left(B_\Omega(o, R')\right) \leq 2 \cdot \frac{\log\left(N_{\mathrm{f}}(\varepsilon, \Omega) p_d(R)\right)}{-\log \varepsilon'}.$$

The lower bound is proven by modifying a method used in [6]. The proof for this considers a convex body Ω in canonical form and uses the upper bound on the number of caps $|C|$ whose Macbeath regions are disjoint and centered at the centroids of the caps' bases. This bound is expressed in terms of the Holmes-Thompson volume of an asymptotic ball centered at an apex o with radius R.

$$|C| = O(Vol^H(AsB(o, R)))$$

The caps serve as *collectors* in the ϵ-approximation of Ω where the complexity above is used to bound the number of collectors produced by a modified method of Lemma 3.2 in [6] such that the convex hull of the set of points S is the approximating polytope. With the relation that the least number of maximal flags is at most a fixed multiple of the number of collectors and the number of flags in the approximating polytope is $O(|C|)$, Vernicos and Walsh show that $N_{\mathrm{f}}(\varepsilon, \Omega)$ is also bounded by $O(Vol^H(AsB(o, R)))$. Finally, observing this bound in the flag-approximability definition and taking the limit as ϵ approaches 0, yields the conclusion $a_f(\Omega) \leq \frac{1}{2}Ent(\Omega)$.

4 Clustering

Clustering is an unsupervised learning task used in machine learning to partition points into meaningful sets called *clusters* based on some measure of similarity. Traditional methods, which consider similarity measures like Euclidean distance and KL-divergence have proven to be effective across many tasks [16]. However, certain applications, like those found in data analysis, text mining, and computer vision, involve data that naturally reside in bounded convex regions whose structure is not accounted for by these measures. One such example arises when clustering categorical distributions. These distributions are typically represented as normalized histograms, but can also be viewed as weighted point set encoding multinoulli distributions in a probability simplex. Prior work attempting to model these distributions in the probability simplex used differential structures by either setting the Riemannian metric tensor to the corresponding Fisher information matrix or defining a KL-divergence induced dual information geometric structure. These approaches have several drawbacks, such as requiring differentiable probability density functions and lacking complete metric properties. Nelson and Sun in [25] address these limitations by introducing a computationally-friendly framework for modeling the geometry of the probability simplex that uses the Hilbert Metric. This section of the survey is primarily based on the work of Nelson and Sun from their 2021 paper *"Clustering in the Hilbert Simplex Geometry"*.

Multinomial Distributions and the Probability Simplex. Multinomial distributions model the probabilities of a set of $d+1$ outcomes in a sequence of m independent events given the probability λ_p^i ($\sum_{i=0}^{d} \lambda_p^i = 1$) for an outcome e_i, $i \in \{0, ..., d\}$. Multinoulli or *categorical* distributions, which are the focus of the work, is a specific case of multinomial distributions, where $m = 1$ and $d > 1$. These distributions serve as the feature representation of individual data points from categorical data and are often expressed as normalized histograms. They can also be viewed as a point $p \in \Delta^d$ on the open probability simplex Δ^d where coordinates reflect the probability of different outcomes. The simplex is embedded in $\mathbb{R}^{d+1}$ on a hyperplane containing points whose coordinates sum up to 1. By interpreting categorical distributions this way, the clustering task can be formulated as follows. Given a set of n points (categorical distributions) $\Lambda = \{p_1, ..., p_n\}$ lying on an open probability simplex Δ^d, we aim to partition Λ into k distinct clusters. Common approaches for this problem use either k-center or center

based k-means++ clustering algorithms with a suitable dissimilarity measure to distinguish between distributions.

Hilbert Simplex Geometry. Since the bounded convex domain Ω required for Hilbert geometry does not need to be a unit ball or smooth at the boundary, we are able to endow bounded polytopes and particularly a d-dimensional open standard simplex Δ^d with a Hilbert geometry. The latter denotes the *Hilbert simplex geometry*, in which the distance between two multinomial distributions p and q is calculated by taking the intersection points t_0 and t_1 of the line $(1-t)p+tq$ with the boundary $\partial\Delta^d$ and computing the following cross ratio:

$$\rho_{\mathrm{HG}}(p,q) = \left|\log \frac{(1-t_0)t_1}{(-t_0)(t_1-1)}\right|.$$

An equivalent formulation on a $(d-1)$-dimensional simplex that uses Birkhoff geometry and does not require finding intersection points (seen below), considers ratios of probability masses assigned to outcomes e_i by distributions p and q.

$$\rho_{\mathrm{HG}}(p,q) = \log \frac{\max_{i\in\{1,\ldots,d\}} \frac{p_i}{q_i}}{\min_{j\in\{1,\ldots,d\}} \frac{p_j}{q_j}} \tag{1}$$

We now note some interesting properties of the Hilbert Simplex Geometry.

1) Due to the cross ratio, by construction, the Hilbert metric is *invariant* to collineations. That is, applying a collineation H on a simplex preserves the Hilbert distance between any two points in the deformed simplex.

$$\rho_{\mathrm{HG}}^{\Delta}(p,q) = \rho_{\mathrm{HG}}^{H(\Delta)}(H(p),H(q)).$$

2) Balls in the Hilbert geometry of a d-dimensional simplex are Euclidean polytopes with $d(d+1)$ facets whose shape stays the same with varying radius and center. They remain as eulidean polytopes at infinitesimal radii, hence the Hilbert simplex geometry is *non-Riemannian*. Note that Hilbert geometry is Finslerian and becomes Riemannian only when the boundary of the convex domain is an ellipsoid.

3) The Hilbert simplex metric space (Δ^d, ρ_{HG}) is *isometric* to a normed metric space $(V^d, ||\cdot||_{NH})$ where $||\cdot||_{NH}$ is a polytope norm. The transformations mapping a point $p=(\lambda^0,\ldots,\lambda^d)$ on the simplex to a point $v(x)=(v^0,\ldots,v^d)\in V^d$ and back are as follows,

$$v^i = \frac{1}{d+1}\left(d\log\lambda^i - \sum_{j\neq i}\log\lambda^j\right).$$

$$\lambda^i = \frac{\exp(v^i)}{\sum_j \exp(v^j)}.$$

where V^d is a unit ball residing in $\mathbb{R}^{d+1}$. Finding the smallest enclosing ball in this vector space with respect to the polytope norm is an LP-type problem and can be solved using Linear programming.

4) The Hilbert metric in the simplex is non-separable and satisfies *information monotonicity*.

5) The Hilbert distance on the simplex can be computed in linear time with a simple, closed formula by considering an equivalent variation norm on a Hilbert projective geometry defined on a cone over the positive orthant.

Comparing Models of the Probability Simplex. Including the Hilbert simplex geometry, Nelson and Sun compare three prominent geometric modelings of the probability simplex. The Fisher-Hotelling Rao geometry models Δ^d with a Riemannian manifold embedded in the positive orthant of a Euclidean d-sphere in $\mathbb{R}^{d+1}$ with a metric tensor derived from the Fisher information matrix. This model has constant positive curvature and demonstrates invariance to reparameterization. Information geometry models Δ^d by a dual structure induced by KL-divergence dissimilarity with dually flat curvature. Unlike FHR, KL-divergence is not a complete metric is it is asymmetric, satisfies information monotonicty, and supports straight line geodesics. By comparison, the Hilbert simplex geometry is a non-Riemannian, complete metric that satisfies information monotonicity and has straight line geodesics (although not unique). It is invariant under projective transformations and does not rely on a manifold representation of the simplex. Curvature in this geometry is negative.

K-means++ Clustering. The widely studied k–means++ algorithm augments the standard k-means algorithm by introducing a random seeding process [4]. Initial cluster centers are chosen arbitrarily or uniformly at random from the set of points with subsequent cluster centers being D^2 weighted. A new center c_i is then chosen with a probability that is proportional to the squared distance D^2 from c_i to the closest center already chosen. In the context of clustering multinomial distributions, the k-means objective of minimizing intra-cluster variances [25] is,

$$E_D(\Lambda, C) = \frac{1}{n}\sum_{i=1}^{n} \min_{j\in\{1,\ldots,k\}} D(p_i : c_j).$$

where Λ is the set of points, C is a set of k cluster centers, and $D(\cdot:\cdot)$ is a dissimilarity measure. The seeding process is adapted by picking an initial center c_1 uniformly at random with subsequent centers or *seeds* $c_2, \ldots, c_k$ selected according to a weighted probability distribution with respect to the chosen divergence.

$$\Pr(c_j = p_i) = \frac{D(p_i, \{c_1, \ldots, c_{j-1}\})}{\sum_{i=1}^{n} D(p_i, \{c_1, \ldots, c_{j-1}\})} \quad (2 \le j \le k).$$

The Hilbert simplex geometry is applied to this algorithm by defining the squared Hilbert distance for D denoted as $D(p, c) = \rho^2_{HG}(p, c)$ where p is a candidate point and c is an already chosen center. Nelson and Sun show that using the Hilbert geometry for k-means++ seeding yields a competitive guarantee of $16(2+\log k)$. The proof for this connects the isometry of the Hilbert simplex geometry to a normed vector space to the general performance guarantee of k-means++ algorithm for any dissimilarity measure. The general performance theorem is used to first establish a competitive bound on the performance of the k-means++ seeding relative to the best possible k-means clustering. This bound depends on two constants κ_1 and κ_2 which are used to evaluate how "metric-like" a given dissimilarity measure is by quantifying how much it deviates from the quasi-triangular inequality and the symmetry inequality respectively. When applying metric properties to this theorem, it shows that any squared metric distance satisfies a 2-approximate triangle inequality resulting in a $16(2 + \log k)$-competitive bound in any metric space with respect to the corresponding squared *metric* distance. A similar argument is made to establish the same competitive bound for any normed vector space using the squared norm as the dissimilarity measure. Finally, because the Hilbert simplex metric space is isometric to a normed vector space equipped with a polytope Hilbert norm, the $16(2+\log k)$-competitive bound directly applies to the squared Hilbert distance $\rho^2_{HG}(p, c)$.

k-center Clustering. Unlike in k-means++ clustering where we try to minimize intra-cluster variances, in k-center clustering, we aim to minimize the maximum dissimilarity between any point and its nearest cluster center with respect to a divergence. In other words, we want to minimize the following cost:

$$f_D(\Lambda, C) = \max_{p_i \in \Lambda} \min_{c_j \in C} D(p_i : c_j).$$

The algorithm starts by randomly picking k cluster centers using the seeding technique from k-means++. Each point on the simplex is then assigned to the closest center, denoted as the label, based on the dissimilarity measure. Then, the cluster centers are updated to minimize the maximum dissimilarity between points in their respective cluster based on the labels, and the distance. Since solving this problem is NP-hard, practical algorithms frequently rely on greedy heuristics for an approximation. A common heuristic, the farthest first travel heuristic, for any metric distance provides an approximation guarantee factor of 2 [25].

Computing the Minimax Center in the Hilbert Simplex Geometry. In the Hilbert simplex geometry, minimizing the maximum dissimilarity between any point and its corresponding cluster center amounts to finding the minimax center (or 1-center) or equivalently the Minimum Enclosing Ball (MEB). Nelson and Sun propose an exact and an approximation algorithm to find the MEB in the simplex. The exact algorithm solves the MEB problem in the isometric normed vector space V^d using an LP-type randomized algorithm and while exact, this algorithm does not scale well in high dimensions. In contrast, the approximation algorithm scales efficiently with the simplex dimension. This algorithm starts by randomly selecting an initial center from the set of points and then iteratively updates the current center c_t by stepping along the geodesic between the farthest point and the previous center c_{t-1} based on the Hilbert metric. This approximation is also referred to as the *geodesic bisection approximation heuristic.*

Performance. Nelson and Sun evaluate the performance of both the k-means++ and k-center algorithms with respect to different dissimilarity measures including FHR, KL-divergence, Hilbert simplex geometry, L2-norm, and L1-norm. Experiments conducted across varying k values, sample sizes, and noise levels demonstrate the superiority of the Hilbert metric over the Riemannian geometries in high dimensions and large noise levels due to the compactness of Hilbert metric balls and a closer reflection of the intrinsic geometry of high-dimensional simplices.

5 Algorithms and Theory

Recently, there has been a growing interest in extending classical computational geometry structures such as Voronoi diagrams and Delaunay Triangulations to the Hilbert geometry. However, translating their Euclidean counterparts to this geometry is non-trivial. For example, in the Hilbert metric space, Delaunay triangulations do not always cover the convex hull of the point set [19] and bisectors separating Voronoi cells are piecewise curves (except in the simplex) as opposed to straight line segments [20]. As such, popular algorithms like Fortune's sweepline algorithm [32] which relies on storing circle events and convex-hull based algorithms like Barber's quickhull algorithm [11] can not be applied in this setting as Hilbert metric balls are Euclidean polytopes and the analogous Hilbert Hull is not necessarily convex.

This section of the survey is primarily based on the works of Gezalyan et al. from their 2021 paper *"Voronoi Diagrams in the Hilbert Metric"*, Gezalyan et al. from their 2023 paper

on *"Delaunay Triangulations in the Hilbert metric"*, and Bumpus et al. from their 2024 paper on *"Analysis of Dynamic Voronoi diagrams in the Hilbert Metric"*.

5.1 Voronoi Diagrams

The Voronoi diagram of a set of points residing on a plane divides the plane into regions or *cells* according to a nearest-neighbor rule, where every point is associated with a region closest to it [9]. Voronoi diagram have been applied in various areas of computer science including computer graphics [39], pattern recognition [18], robotics [28, 40], and operations research. In particular, four of the most common uses motivating this structure are highlighted by Aurenhammer in his 1991 survey which include associative file searching, cluster analysis, scheduling record accesses, and collision detection. We now note the general Euclidean characterization of Voronoi diagrams. Let S denote a set of n points, also referred to as sites, which reside in a plane. The Voronoi *cell* $V_S(p)$ for any point $p \in S$ is defined as a subset of the plane that is at least as close to p as to any other point p',

$$V_S(p) = \left\{ q \in \mathbb{R}^2 : d_E(q,p) \leq d_E(q,p'),\ \forall p' \in S \setminus \{p\} \right\}.$$

where $d_E(\cdot,\cdot)$ is the Euclidean distance. Each cell can be viewed as the intersection of $n-1$ half-planes. The Voronoi diagram of S is then the resulting polygonal partition of the plane induced by the Voronoi cell of each point $p \in S$.

Characterization of Voronoi Diagrams in the Hilbert metric space In contrast to the Euclidean characterization, Voronoi diagrams in the Hilbert metric space, named *Hilbert Voronoi diagrams*, are defined with respect to a convex polygon K residing in $\mathbb{R}^2$. The set S now lies within the interior of K. The Voronoi *cell* $V_S(p)$ for any point $p \in S$ is defined as

$$V_S(p) = \left\{ q \in K, K \in \mathbb{R}^2 : d_H(q,p) \leq d_H(q,p'),\ \forall p' \in S \setminus \{p\} \right\}.$$

where d_H is the Hilbert distance. The Hilbert Voronoi diagram of S, $\mathrm{Vor}_K(S)$, is then defined as the cell complex of K constructed from the Voronoi cells $V(p)$ for all $p \in S$. We note two interesting properties of Hilbert Voronoi diagrams.

▶ **Observation 2** (Star Shaped). *With respect to the associated site, Voronoi cells in the Hilbert metric are stars*

▶ **Observation 3** (2-dimensional Bisectors). *When the line passing through any pair of sites and the lines extending any two edges of K are not conincident at a common point, the bisectors separating the Voronoi cells become 2-dimensional.*

Hilbert Metric Bisectors. The Hilbert bisector of two points p and p' where $p, p' \in \mathrm{int}(K)$, is defined as the set of points $z \in K$ such that $d_H(z,p) = d_H(z,p')$. Unlike Euclidean bisectors which are straight lines, Hilbert bisectors are generally piecewise curves which depend on the boundary edges of K. To better understand this piecewise nature, Bumpus et al [13] define *spokes* and *sectors*.

▶ **Definition 4** (Chord). *The chord $\chi(a,b)$ of two points a and b is the intersection of the line passing through a and b and ∂K*

▶ **Definition 5** (Spokes). *A spoke, with respect to a vertex v of K is the intersection between the line passing through a site $s \in K$ and v and ∂K*

The spokes of sites p and p' subdivide the interior of the K into polygonal regions called *sectors*.

▶ **Definition 6** (Sector). *A sector is defined by four boundary edges which are not necessarily distinct. These edges are obtained from chords $\chi(a, p)$ and $\chi(a, p')$ where a is any point within the sector.*

The Hilbert bisector of two points p and p' is made up of $O(m)$ *bisector segments*, where m is the total number of sides of K. Each segment is formed by a set of points lying in an individual sector s and can be described by a conic section whose coefficients depend on the line equations of the four boundary edges corresponding to s [13]. Breakpoints or *joints* occur when we travel along the bisector from the boundary of one sector to another [20]. Using these segments, Gezalyan and Mount show that the Hilbert Voronoi diagram has a combinatorial complexity of $O(mn)$. To compute this diagram they propose a randomized incremental and a divide and conquer algorithm.

Randomized Incremental Algorithm. Given n sites in an m-sided polygon K, the randomized incremental algorithm for constructing the Hilbert Voronoi diagram randomly permutes the sites and adds each site to the diagram one by one, similar to common incremental approaches for constructing Euclidean Voronoi diagrams. This algorithm also produces a point-location data structure which queries a given point and returns the closest site location in an expected time of $O(\log n \log mn)$. We now detail constructing, updating, and querying this structure.

Construction: While the diagram is being constructed, line segment spokes between the current site s and the boundary points of $V(s)$ are added to a planar subdivision data structure (i.e doubly-edged linked list). These points include:

1. Vertices of $V(s)$
2. Joints lying along edges of $V(s)$
3. Intersection points of an edge of $V(s)$ and ∂K
4. Vertices of K that happen to lie in $V(s)$

This construction results in a *triangulation* whose size is proportional to the size of diagram.

Update: Gezalyan and Mount leverage a *history-DAG* which encodes a record of structural changes made to the diagram. Leaf nodes in this graph represent individual triangles in the triangulation and the root node represents the entire polygon K. When a new site is inserted, the old triangles in the history-DAG are "destroyed" by storing a reference to the new triangle that replaced it. The new triangle corresponds to the inserted site that caused the destruction.

Query: Given a query point q, point location on the history-DAG is performed. Let p_j denote the site corresponding to q's triangle before p_i the point that caused the change was inserted. There are two possible cases, either the query is still closest to the old site or its now closer to the new site. Compute the Hilbert distance between the query and these sites to determine the closer one. Then, perform a radial binary search around this site to figure out the current triangle containing q.

Divide and Conquer Algorithm. The divide and conquer algorithm constructs the Hilbert Voronoi diagram for n sites contained in an m-sided polygon in $O(nm \log n)$ time and appears similar to the Euclidean counterpart. The algorithm starts by partitioning the

sites into two subsets of roughly equal size with respect to a vertical line. Hilbert Voronoi diagrams are computed recursively on these subsets. The merge step involves computing the bisector between the subsets with a single bottom-up traversal in $O(mn)$ time. This traversal constructs bisector *components*, a collection of curves, which are formed by the concatenation of Voronoi edges of a site from the left subset and a site from the right subset. After constructing the bisector, portions of the left subset's Voronoi diagram that lie on the right side of the bisector and portions of the right subset's Voronoi diagram that lie on the left side are removed.

5.2 Delaunay Triangulations

The Delaunay Triangulation of a set of points P is a planar subdivision whose bounded faces are triangles with vertices drawn from points in P [14] and can be described as the dual graph of the Voronoi diagram of P. Delaunay triangulations have been applied in many areas including computational geometry, multi-agent systems, and spatial analysis. In particular, an iterative optimization algorithm called LLoyd's algorithm, leverages Delaunay triangulations in finding subsets of the Euclidean space with an evenly spaced set of points and partitioning them into well-shaped convex cells. Doing so refines the quality of the resulting mesh by reducing the occurrences of overly acute triangles which tend to cause numerical instability [15]. In spatial clustering, this structure is used as part of an adaptive spatial clustering algorithm called ASCDT where spatial clusters with complex shapes and non-uniform densities can be identified using a proximity definition based on the triangulation and statistical features of its edges [15].

In the Euclidean setting, common algorithms for constructing Delaunay triangulations such as flip and incremental search algorithms rely on *in-circle tests* which determine whether a point lies inside a triangle's circumcircle. The flip algorithm, for example, iteratively inspects pairs triangles for violations of the *empty circumcircle condition*. That is, in any valid Delauanay triangulation, the circumcircle of any face must not contain a point from P in its interior. If a violation is detected, the shared edge is flipped and this continues until convergence. Similarly, in the incremental construction algorithm, points are added one by one to some existing triangle. This triangle is then identified and corrections are performed to preserve the *Delaunay property* (empty circumcircle condition) by flipping edges. This way, the Delaunay property is maintained at the end of each insertion and consequently as a whole by the end of the algorithm. Similar to the case of Hilbert Voronoi diagrams, these algorithms can not be directly applied in the Hilbert geometry as the nature of circumcircles and the conditions for satisfying the *empty Hilbert circumcircle condition* are more complex than the Euclidean case. For example, in the Hilbert geometry, it can not be guaranteed that any three non-collinear points form a unique circumcircle and as such there can exist triangles that may not admit any Hilbert circumcircle.

Characterizing the Hilbert Delaunay Triangulation The *Hilbert Delaunay Triangulation* of a set of sites P, $\mathrm{DT}_\Omega(P)$, is described by the dual of the Hilbert Voronoi diagram. That is, any two sites $p, q \in P$ are connected by an edge in the triangulation if $\mathrm{Vor}_\Omega(p)$ and $\mathrm{Vor}_\Omega(q)$ are adjacent. This triangulation adapts common concepts from the Euclidean counterpart using Hilbert balls such as two sites being adjacent if and only if they lie on the boundary of a Hilbert ball which has no other sites in its interior, and three points form a triangle if and only if they lie on the boundary of some Hilbert ball. In the latter case, the Hilbert ball is said to *circumscribe* the triangle. Gezalyan et al. detail several properties of the Hilbert Delaunay Triangulation:

1. Any triangle in $int(\Omega)$ admits at most one Hilbert circumcircle
2. If P is discrete then $\mathrm{DT}_\Omega(P)$ is a *planar* graph spanning P
3. $\mathrm{MST}_\Omega(P) \subseteq \mathrm{RNG}_\Omega(P) \subseteq \mathrm{DT}_\Omega(P)$, where $\mathrm{MST}_\Omega(P)$ is the Hilbert Minimum Spanning Tree of P in which the edge weights are Hilbert distances and $\mathrm{RNG}_\Omega(P)$ is the Hilbert Relative Neighborhood Graph of P
4. $\mathrm{DT}_\Omega(P)$ does not necessarily cover the convex hull of P

Hilbert Circumcircles. The algorithm proposed by Gezalyan et al. to construct the Hilbert Delauanay triangulation for a set of points adapts the well known randomized incremental algorithm for Euclidean Delauanay triangulations [21]. This algorithm relies on an *incircle* test which determines whether a newly inserted site lies on the interior of a circumcircle of a triangle formed by three existing sites. This is also referred to as the *empty circumcircle condition* and understanding the Hilbert metric analog of this condition is important. However, this is non-trivial as the Euclidean property that any triple of non-collinear points must lie on the boundary of a unique Euclidean ball does not extend to the Hilbert geometry. Gezalyan et al. explore the conditions when a triangle admits a Hilbert circumcircle and how to compute Hilbert circumcircles.

To determine the conditions when a triangle admits the Hilbert circumcircle, Gezalyan et al. characterize *balls at infinity* in the Hilbert metric. The centers of these balls are located on the boundary of Ω and approach a convex polygon lying on the interior of Ω with respect to $\partial\Omega$ as the limit. Using this characterization, Gezalyan et al. define overlap and outer regions for any two points $p, q \in \Omega$. The overlap region is the intersection of the balls at infinity centered at the endpoints of the (p, q)-bisector (Hilbert) intersecting $\partial\Omega$. The outer region is Ω excluding the union of these balls. With these regions, the condition for a triangle Δpqr to admit a Hilbert circumcircle is if and only if the site r does not belong in the union of the overlap and outer regions of p and q.

Given three non-collinear sites p, q, and r, computing the Hilbert circumcircle of triangle Δpqr amounts to finding a point c equidistant from the points. Equivalently, this means computing the point of intersection of the (p, q)- and (q, r)-bisectors. The algorithm for computing the circumcircle proceeds by determining whether the (p, q)- and (q, r)-bisectors intersect within the interior of Ω. This is done by computing the endpoints of the bisectors in $O(\log^2 m)$ time by performing a binary search routine on $\partial\Omega$ and then checking if they alternate between pairs (p, q) and (q, r) along $\partial\Omega$. The algorithm then finds the unique circumcenter c by performing a binary search along an angular interval from rays $\overrightarrow{pv_q}$ and $\overrightarrow{pv_r}$ where v_q and v_r are bisector endpoints. In $O(\log^3 m)$ time, this search narrows the interval to three double wedge regions defined by consecutive spokes corresponding to each site, whose intersection contains bisector segments described by simple conics. Finally, the algorithm formulates these conics and determines their intersection point in constant time.

Randomized Incremental Algorithm. To account for the fact that the Hilbert Delaunay triangulation does not necessarily cover the convex convex hull of P, Gezalyan et al. augment the triangulation with additional components for convenient construction. These components include:

1. **Tooth:** The triangle Δpqx where for any two sites $p, q \in \mathrm{int}(\Omega)$, whose edge corresponds to an external face of the triangulation lying to the left, the (p, q)-bisector has an endpoint x intersecting $\partial\Omega$.
2. **Gap:** Regions of Ω that lie outside the standard triangles (formed by any three sites in P) and teeth. Each gap has an associated site, sides of two teeth, and a convex, polygonal

chain of line segments.

Given a set of sites P lying within Ω, the algorithm starts by randomly permuting the sites. Then, it inserts two sites a and b and creates an initial augmented triangulation by computing the endpoints x and y of the (a, b)-bisector and adding the edges ab, ax, ay, bx, by to the triangulation. ax, ay, bx, by describe the two teeth lying on opposite sides of the edge ab. After this, the algorithm incrementally inserts the remaining points, updating the triangulation along the way. The insertion routine for the remaining points works as follows. There are three cases, respective of the component into which the inserted site falls into. In the case of a standard triangle, the new site is connected to vertices of the triangle. In the case of a tooth, the boundary vertex connecting the sites part of the tooth and the new site is removed and two new teeth are created based on the bisectors of the three sites (two existing and one new). Finally, in the case of a gap, the new and existing sites are connected and two new teeth are created with respect to the bisectors of the new site.

Participating in the insertion routine runs the risk of not satisfying the Delaunay empty circumcircle condition, producing a geometrically invalid triangulation. This is due to the additional components added to the triangulation. To ensure a topologically correct triangulation, Gezalyan et al. propose two correcting procedures FLIPEDGE and FLIPTOOTH. FLIPEDGE is applied to the edges of newly formed standard triangles where if a new site fails the incircle test of the triangle Δabc neighoring the triangle Δpab where p lies to left of the directed edge $\overrightarrow{ab}$, then the edge $\overrightarrow{ab}$ is removed. When is Δabc is standard, the edge pc is added and the edges ac and cb are processed. If Δabc is not standard, two new teeth based on the $(pa)-$ and $(bp)-$bisectors are created. FIXTOOTH is applied to teeth in situations where a newly created boundary vertex causes overlapping behavior between two or more teeth. The procedure checks for potential overlaps between neighboring teeth and if detected, the overlapping teeth are replaced by a standard triangle and a tooth. Then, FLIPEDGE and FIXTOOTH are called on an edge of the new triangle and tooth respectively.

Gezalyan et al. show that this randomized incremental algorithm, with the augmented triangulation and corrections runs in expected time $O(n(\log n + \log^3 m))$ for n sites in an m-sided polygon.

Hilbert Hull. The *Hilbert Hull* of a set of points P is the region covered by triangles of the Hilbert Delaunay triangulation of P. Gezalyan et al. present an $O(nh \log^2 m)$ algorithm based on the Jarvis March algorithm to construct the Hilbert hull of a set of points where h is the number of edges on the boundary of the hull. The algorithm starts by finding a point $p_0 \in P$ lying on the boundary of an empty ball at infinity centered at a point $x_0 \in \partial\Omega$. Given the initial pair is (p_0, x_0), subsequent pairs (p_i, x_i) are obtained by computing the bisector endpoints of all points in P excluding p_{i-1}. The next point on the hull is selected from the smallest point derived from a lexicographical ordering induced by (p_{i-1}, x_{i-1}). Pairs are computed until the starting pair is reached or passed.

6 Quantum Information Theory

The application of the Hilbert metric in quantum information theory is enabled by the presence of convex cones representing quantum operators including the set of positive semi-definite operators and the subset of separable operators. Important concepts that arise in the discussion of these cones are certain classes of cone-preserving linear maps and local operations with classical communication (LOCC) maps [31]. In this work, finite-dimensional real vector spaces $\mathcal{V}$ are thought of as the space of Hermitian matrices $\mathcal{M}_d(\mathbb{C})$, thus $\mathcal{V}$

is isomorphic to $\mathbb{R}^{d^2}$. Convex cones are then defined as subsets of this space for which $\alpha\mathcal{C} + \beta\mathcal{C} \subseteq \mathcal{C}$ where α and β non-negative constants. Reed et al. focus specifically on *proper* cones which represent convex sets of interest such as density matrices, separable states, PPT states, and more. A proper cone is a convex cone that is closed, pointed and solid, where

1. Pointed: The only element shared between $\mathcal{C}$ and its reflection $-\mathcal{C}$ is the zero vector
2. Solid: The span of $\mathcal{C}$ is the vector space $\mathcal{V}$

These cones correspond to *partial orders* in $\mathcal{V}$ where $a \geq b$ refers to vectors $a - b \in C$. This ordering helps to define positive elements in the vector space which is useful when describing convex sets such as the space of positive semi-definite matrices in which operators are ordered as $a \geq b$. The main proper cones discussed in the work involving the Hilbert metric can be divded into two areas:

$\mathcal{C} = S_+$	Space of positive semi-definite matrices
$\mathcal{C} = S_{\text{PPT}}$	Set of positive partial transpose (PPT) matrices
$\mathcal{C} = S_{\text{PPT}+}$	$S_+ \cap\ S_{\text{PPT}}$ generated by all PPT states
$\mathcal{C} = S_{\text{SEP}}$	Set of separable matrices
$\mathcal{C} = (\mathbb{R}_+)^d$	Vectors with non-negative entries (Perron-Frobenius Theory)

Table 1 Base Norms & Negativities

$\mathcal{C}_{M_+} = S_+$	Set of POVM elements
$\mathcal{C}_{M_{LO}}$	Set of 2-outcome measurements with LO & no communication
$\mathcal{C}_{M_{LOCC}}$	Set of 2-outcome measurements implementable by LOCC
$\mathcal{C}_{M_{SEP}}$	All separable measurements on an n-partite quantum system
$\mathcal{C}_{M_{PPT+}}$	All PPT measurements

Table 2 Distinguishability Measures

This section of the survey is primarily based on the work of Reed et al. from their 2021 paper *"Hilbert Projective Metric in Quantum Information theory"*.

Hilbert Projective Metric in Quantum Information Theory In a proper cone $\mathcal{C}$ for any two elements $a, b \in \mathcal{C}\ \{0\}$ the definition of the *Hilbert Projective Metric* mirrors the formulation seen in [23],

$$h(a, b) = \ln(\sup(a/b)\sup(b/a))$$

where $\sup(a/b)$ is defined as,

$$\sup(a/b) = \inf\{\lambda \in \mathbb{R}\ |a \leq_{\mathcal{C}} \lambda b\}$$

Recalling partial order, λ in this definition provides the scaling factors for b such that it becomes at least smaller or larger than a. This definition depends only on the direction of a and b and is a *projective* metric with respect to $\mathcal{C}$. In this vein, an application discussed by Reeb et al. considers the cone of positive semi-definite matrices S_+ containing all density matrices of a d-dimensional quantum system and max-relative entropy. Specifically, for normalized quantum states ρ and σ, $\sup_{S_+}(\rho, \sigma)$ represents the largest possible probability ratio of any measurement outcome on ρ compared to σ, equaling the max-relative entropy of

ρ and σ up to a logarithm factor [31].

Contraction Properties for Cone-preserving maps. Some of the main applications of the Hilbert Projective metric is in the context of *cone-preserving* maps, where a transformation T sends one cone $\mathcal{C}$, corresponding to a set of positive elements to another cone $\mathcal{C}'$. These cones represent partially ordered vector spaces $\mathcal{V}$ and $\mathcal{V}'$ respectively. The *projective diameter* $\Delta(T)$, which represents the largest possible Hilbert distance between any two points in $T(C) \subseteq C'$, can then be formulated as,

$$\Delta(T) = \sup_{a,b \in \mathcal{C} \setminus \{0\}} h_{\mathcal{C}'}(T(a), T(b)).$$

The Hilbert metric serves as a convenient theoretical tool [31] for the analysis of contraction bounds and based on the above formulation Reeb et al. discuss several contraction properties. The Birkhoff-Hoph theorem states that any cone-preserving map T acts as a *contraction* with respect to the Hilbert metric and $\tanh[\Delta(T)/4]$ is the best possible contraction coefficient [31]. The contraction coefficient provides a spectral bound on the projective diameter and is bounded by a contraction ratio of T with respect to any projective metric D. The later refers to the *uniqueness of the Hilbert metric* where for any strictly contracting linear map sending a proper cone to its interior with respect to D, there exists a continuous and strictly increasing function such that $D(a, b) = h_{\mathcal{C}}(a, b), \forall a, b \in \text{int}(\mathcal{C})$. The bound in particular is,

$$\tanh\left(\frac{\Delta(T)}{4}\right) \leq \sup_{a,b \in \mathcal{C} \setminus \{0\}} \left\{ \frac{D(T(a), T(b))}{D(a, b)} \;\middle|\; D(a, b) > 0 \right\}$$

Additional contraction properties relate base norms and negativities to the projective diameter of cone-preserving maps. The *base* $\mathcal{B}$ of a proper cone $\mathcal{C} \subset \mathcal{V}$ is a convex subset of $\mathcal{C}$ where any element $c \in \mathcal{C} \setminus \{0\}$ lies on a ray intersecting at exactly one point $b \in \mathcal{B}$. A base of a proper cone equips the corresponding vector space $\mathcal{V}$ with a norm referred to as the *base-norm* [31]. *Negativity* is a value $\mathcal{N}$ that represents how far a quantum state v is from the cone where $\mathcal{N}(v) = 0 \iff v \in \mathcal{C}$ and is used to quantify entanglement in a bipartite quantum system.

A common inequality seen in quantum information theory called the Ruskai's trace-norm contraction inequality which states that the trace-norm distance of two quantum states does not increase by applying a quantum channel. In other words, the states do not become more distinguishable when acted upon by a channel. Formally, given quantum states ρ_1, ρ_2 and a channel T, then

$$\|T(\rho_1) - T(\rho_2)\|_1 \leq \|\rho_1 - \rho_2\|_1$$

Reeb et al. show that leveraging the Hilbert projective metric, one can multiply the trace-norm of ρ_1 and ρ_2 by the contraction coefficient which improves this bound, yielding a stricter contraction.

An alternative interpretation of the trace norm is observed in the discussion of distinguishability measures in which the distance quantifies the best possible distinguishability between two quantum states across all physical measurements [31]. Reeb et al. show that that this distance or the *maximal bias* in distinguishing the states can be upper bounded by the Hilbert projective geometry. In particular, in the space of implementable measurements $\mathcal{C}_{S+}$, the trace distance between two quantum states ρ_1 and ρ_2 belonging to the base of $\mathcal{C}_{S+}$ can be bounded as follows,

$$\frac{1}{2}\|\rho_1 - \rho_2\|_1 \leq \tanh\left(\frac{h_{S_+}(\rho_1, \rho_2)}{4}\right).$$

Reeb et al. also provide bounds on two other popular distinguishability measures including *fidelity*, which generalizes the overlap between mixed and pure states, and *Chernoff bound* which is the asymptotic rate at which the symmetric quantum hypothesis testing error reduces. Formally, given two density matrices $\rho_1, \rho_2 \in \mathcal{M}_d(\mathbb{C})$ and the cone $\mathcal{C}_{S+}$, fidelity F is bounded as,

$$\sqrt{1 - F(\rho_1, \rho_2)^2} \leq \tanh\left(\frac{h_{S_+}(\rho_1, \rho_2)}{4}\right),$$

and the Chernoff bound is bounded as,

$$-\ln \min_{0 \leq s \leq 1} \operatorname{tr}\left[\rho_1^s \rho_2^{1-s}\right] \leq \frac{h_{S_+}(\rho_1, \rho_2)}{2}.$$

Applications in Entanglement Theory. In entanglement theory, the Hilbert projective metric is used to specify a non-trivial contraction ratio for quantifying the decrease in entanglement described by entanglement monotones like negativities. Negativity contraction normally denotes that on average, entanglement is non-increasing under the application of LOCC operations. That is, if an LOCC operation maps a quantum state ρ to a state ρ_i with some probability p_i, then the decrease in entanglement is upper bounded by the initial negativity of ρ. Formally, this means that $\eta = 1$ and

$$\sum_i p_i \mathcal{N}(\rho_i) \leq \eta\, \mathcal{N}(\rho)$$

Applying the Hilbert projective metric shows that the *generalized* negativity is an entanglement monotone by specifying $\eta \leq 1$ in terms of the projective diameter. That is, given a linear, cone-preserving, non-deterministic operation T_i mapping a vector space $\mathcal{V}$ to a vector space $\mathcal{V}_i$, with proper cones $\mathcal{C}$ and $\mathcal{C}_i$ and bases $\mathcal{B}$ and $\mathcal{B}_i$, the negativity contraction is

$$\sum_{i=1}^{N} p_i \mathcal{N}_{\mathcal{B}_i}(\rho_i) \leq \mathcal{N}_{\mathcal{B}}(\rho) \tanh\left(\frac{\max_i \Delta(T_i)}{4}\right)$$

Operational Interpretation of the Hilbert Metric Measures such as the Hilbert metric, fidelity, trace-norm, and χ-divergence all share a contraction property via the Birkhoff-Hopf contraction theorem: when a quantum channel is applied to two quantum states, their distance does not increase [31]. Only from the contractivity of the Hilbert metric however, can we decide the existence of a *completely positive* map. That is, given two pairs of quantum states (ρ_1, ρ_1') and (ρ_2, ρ_2'), a completely positive map $\rho_i \mapsto \rho_i'$ can be realized in terms of a probabilistic quantum operation $T(\rho_i) = p_i \rho_i'$ if and only if their Hilbert distance is non-increasing:

$$h_{\mathcal{S}_+}(\rho_1, \rho_2) \geq h_{\mathcal{S}_+}(\rho_1', \rho_2').$$

In other words, the Hilbert projective metric serves to provide a necessary and sufficient condition for the existence of T [31].

7 Conclusion

In this work, we surveyed several algorithms and applications of the Hilbert metric appearing in computational settings. While this metric has been extensively studied in mathematical

contexts, applying this metric in these settings reveal useful observations such as how its projective invariance and contractivity make it a convenient tool for analyzing contraction bounds in quantum information theory, its superior performance in clustering tasks on a probability simplex compared to other popular measures, and its similarity to Macbeath regions which connected these regions to Delone sets for achieving better query and storage complexities for approximate polytope membership.

Furthermore, we find that the non-Riemannian nature of this geometry leads to unusual properties like bisectors sometimes being 2-dimensional, the Hilbert hull which is Hilbert geometry analog to the convex hull not necessarily being convex, three non-collinear points not always guaranteeing a unique circumcircle, and more. These properties make it harder to adapt traditional algorithms found in Euclidean settings and to perform computations involving the metric (i.e computing the exact projective diameter in all cases). Thus, we hope this survey helps to further the exploration of the Hilbert metric in order to develop more practical algorithms and applications.

8 Software & Code

The table below lists several software and code contributions accompanying the study of the Hilbert metric.

Paper	Link
On the balls of Hilbert Polygonal Geometry [24]	Demo Hilbert Geometry
Ipelets for the Convex Polygonal Geometry [29]	Ipelets for the Convex Polygonal Geometry
Software for the Thompson and Funk Geometries [10]	Funk Geo Visualizer
Dynamic Voronoi Diagrams in the Hilbert Metric [13]	Dynamic Voronoi Diagrams in the Hilbert Metric
Clustering in the Hilbert Simplex Geometry [25]	Clustering in the Hilbert Simplex Geometry

References

1 Ahmed Abdelkader, Sunil Arya, Guilherme D. da Fonseca, and David M. Mount. Approximate nearest neighbor searching with non-euclidean and weighted distances, 2023. URL: `https://arxiv.org/abs/2306.15621`, `arXiv:2306.15621`.

2 Ahmed Abdelkader and David M Mount. Economical delone sets for approximating convex bodies. In *16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

3 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.

4 David Arthur and Sergei Vassilvitskii. k-means++: the advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '07, page 1027–1035, USA, 2007. Society for Industrial and Applied Mathematics.

5 Rahul Arya, Sunil Arya, Guilherme D da Fonseca, and David Mount. Optimal bound on the combinatorial complexity of approximating polytopes. *ACM Transactions on Algorithms*, 18(4):1–29, 2022.

6 Sunil Arya, Guilherme D Da Fonseca, and David M Mount. On the combinatorial complexity of approximating polytopes. *Discrete & Computational Geometry*, 58:849–870, 2017.

7 Sunil Arya, Guilherme D Da Fonseca, and David M Mount. Approximate polytope membership queries. *SIAM Journal on Computing*, 47(1):1–51, 2018.

8 Sunil Arya, Guilherme D. da Fonseca, and David M. Mount. Economical convex coverings and applications. *SIAM Journal on Computing*, 53(4):1002–1038, 2024. arXiv:https://doi.org/10.1137/23M1568351, doi:10.1137/23M1568351.

9 Franz Aurenhammer. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM computing surveys (CSUR)*, 23(3):345–405, 1991.

10 Hridhaan Banerjee, Carmen Isabel Day, Auguste H. Gezalyan, Olga Golovatskaia, Megan Hunleth, Sarah Hwang, Nithin Parepally, Lucy Wang, and David M. Mount. Software for the thompson and funk polygonal geometry, 2025. URL: https://arxiv.org/abs/2503.01988, arXiv:2503.01988.

11 C Bradford Barber, David P Dobkin, and Hannu Huhdanpaa. The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software (TOMS)*, 22(4):469–483, 1996.

12 Michael A. Bender, Martin Farach-Colton, Mayank Goswami, Rob Johnson, Samuel McCauley, and Shikha Singh. Bloom filters, adaptivity, and the dictionary problem. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 182–193, 2018. doi:10.1109/FOCS.2018.00026.

13 Madeline Bumpus, Xufeng Caesar Dai, Auguste H. Gezalyan, Sam Munoz, Renita Santhoshkumar, Songyu Ye, and David M. Mount. Analysis of dynamic voronoi diagrams in the Hilbert metric, 2024. URL: https://arxiv.org/abs/2304.02745, arXiv:2304.02745.

14 Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Cheong Schwarzkopf. *Delaunay Triangulations*, pages 183–210. Springer Berlin Heidelberg, Berlin, Heidelberg, 2000. doi:10.1007/978-3-662-04245-8_9.

15 Yahia S. Elshakhs, Kyriakos M. Deliparaschos, Themistoklis Charalambous, Gabriele Oliva, and Argyrios Zolotas. A comprehensive survey on delaunay triangulation: Applications, algorithms, and implementations over cpus, gpus, and fpgas. *IEEE Access*, 12:12562–12585, 2024. doi:10.1109/ACCESS.2024.3354709.

16 Absalom E Ezugwu, Abiodun M Ikotun, Olaide O Oyelade, Laith Abualigah, Jeffery O Agushaka, Christopher I Eke, and Andronicus A Akinyelu. A comprehensive survey of clustering algorithms: State-of-the-art machine learning applications, taxonomy, challenges, and future research prospects. *Engineering Applications of Artificial Intelligence*, 110:104743, 2022.

17 Dmitry Faifman, Constantin Vernicos, and Cormac Walsh. Volume growth of funk geometry and the flags of polytopes. *arXiv preprint arXiv:2306.09268*, 2023.

18 Camillo Gentile and Mario Sznaier. An improved voronoi-diagram-based neural net for pattern classification. *IEEE Transactions on Neural Networks*, 12(5):1227–1234, 2001.

19 Auguste Gezalyan, Soo Kim, Carlos Lopez, Daniel Skora, Zofia Stefankovic, and David M Mount. Delaunay triangulations in the Hilbert metric. *arXiv preprint arXiv:2312.05987*, 2023.

20 Auguste H Gezalyan and David M Mount. Voronoi diagrams in the Hilbert metric. *arXiv preprint arXiv:2112.03056*, 2021.

21 Leonidas J Guibas, Donald E Knuth, and Micha Sharir. Randomized incremental construction of delaunay and voronoi diagrams. *Algorithmica*, 7(1):381–413, 1992.

22 D. Hilbert. Ueber die gerade Linie als kürzeste Verbindung zweier Punkte. *Math. Annalen*, 46:91–96, 1895.

23 Bas Lemmens and Roger Nussbaum. Birkhoff's version of Hilbert's metric and its applications in analysis. *Handbook of Hilbert geometry*, pages 275–303, 2014.

24 Frank Nielsen and Laetitia Shao. On balls in a Hilbert polygonal geometry (multimedia contribution). In *33rd International Symposium on Computational Geometry (SoCG 2017)*, pages 67–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2017.

25 Frank Nielsen and Ke Sun. Clustering in Hilbert simplex geometry. *CoRR*, abs/1704.00454, 2017. URL: http://arxiv.org/abs/1704.00454, arXiv:1704.00454.

26 Frank Nielsen and Ke Sun. Clustering in Hilbert's projective geometry: The case studies of the probability simplex and the elliptope of correlation matrices. *Geometric structures of information*, pages 297–331, 2019.

27 Frank Nielsen and Ke Sun. Non-linear embeddings in hilbert simplex geometry. In *Topological, Algebraic and Geometric Learning Workshops 2023*, pages 254–266. PMLR, 2023.

28 Melih Özcan and Ulas Yaman. A continuous path planning approach on voronoi diagrams for robotics and manufacturing applications. *Procedia Manufacturing*, 38:1–8, 2019.

29 Nithin Parepally, Ainesh Chatterjee, Auguste H. Gezalyan, Hongyang Du, Sukrit Mangla, Kenny Wu, Sarah Hwang, and David M. Mount. Ipelets for the Convex Polygonal Geometry. In Wolfgang Mulzer and Jeff M. Phillips, editors, *40th International Symposium on Computational Geometry (SoCG 2024)*, volume 293 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 92:1–92:7, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: `https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.SoCG.2024.92`, `doi:10.4230/LIPIcs.SoCG.2024.92`.

30 Nithin Parepally, Ainesh Chatterjee, Auguste H Gezalyan, Hongyang Du, Sukrit Mangla, Kenny Wu, Sarah Hwang, and David M Mount. Ipelets for the convex polygonal geometry (media exposition). In *40th International Symposium on Computational Geometry (SoCG 2024)*, pages 92–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.

31 David Reeb, Michael J Kastoryano, and Michael M Wolf. Hilbert's projective metric in quantum information theory. *Journal of mathematical physics*, 52(8), 2011.

32 Peter Su and Robert L. Scot Drysdale. A comparison of sequential delaunay triangulation algorithms. *Computational Geometry*, 7(5):361–385, 1997. 11th ACM Symposium on Computational Geometry. URL: `https://www.sciencedirect.com/science/article/pii/S0925772196000259`, `doi:10.1016/S0925-7721(96)00025-9`.

33 Elias Szabo-Wexler. Approximate membership of sets: A survey.

34 Julian Vanecek, Auguste H Gezalyan, and David M Mount. Support vector machines in the hilbert geometry.

35 Constantin Vernicos. Asymptotic volume in Hilbert geometries. *Indiana University Mathematics Journal*, pages 1431–1441, 2013.

36 Constantin Vernicos. On the Hilbert geometry of convex polytopes. *Handbook of Hilbert geometry*, 22:111–125, 2014.

37 Constantin Vernicos and Cormac Walsh. Flag-approximability of convex bodies and volume growth of Hilbert geometries. In *Annales Scientifiques de l'École Normale Supérieure*, volume 54, pages 1297–1315, 2021.

38 Xinyue Wei, Minghua Liu, Zhan Ling, and Hao Su. Approximate convex decomposition for 3d meshes with collision-aware concavity and tree search. *ACM Transactions on Graphics (TOG)*, 41(4):1–18, 2022.

39 Dong-Ming Yan, Wenping Wang, Bruno Lévy, and Yang Liu. Efficient computation of clipped voronoi diagram for mesh generation. *Computer-Aided Design*, 45(4):843–852, 2013.

40 Meng Zhou, Jianyu Li, Chang Wang, Jing Wang, and Li Wang. Applications of voronoi diagrams in multi-robot coverage: A review. *Journal of Marine Science and Engineering*, 12(6):1022, 2024.