Economical Delone Sets for Approximating Convex Bodies

Ahmed Abdelkader
Department of Computer Science
University of Maryland, College Park MD, USA
akader@cs.umd.edu
https://orcid.org/0000-0002-6749-1807

David M. Mount
Department of Computer Science and Institute of Advanced Computer Studies
University of Maryland, College Park MD, USA
mount@cs.umd.edu

Abstract

Convex bodies are ubiquitous in computational geometry and optimization theory. The high combinatorial complexity of multidimensional convex polytopes has motivated the development of algorithms and data structures for approximate representations. This paper demonstrates an intriguing connection between convex approximation and the classical concept of Delone sets from the theory of metric spaces. It shows that with the help of a classical structure from convexity theory, called a Macbeath region, it is possible to construct an ε-approximation of any convex body as the union of $O(1/\varepsilon^{(d-1)/2})$ ellipsoids, where the center points of these ellipsoids form a Delone set in the Hilbert metric associated with the convex body. Furthermore, a hierarchy of such approximations yields a data structure that answers ε-approximate polytope membership queries in $O(\log(1/\varepsilon))$ time. This matches the best asymptotic results for this problem, by a data structure that both is simpler and arguably more elegant.

2012 ACM Subject Classification Theory of computation → Randomness, geometry and discrete structures → Computational geometry

Keywords and phrases Approximate polytope membership, Macbeath regions, Delone sets, Hilbert geometry

Digital Object Identifier 10.4230/LIPIcs.SWAT.2018.4

Funding Research supported by NSF grant CCF–1618866.

1 Introduction

We consider the following fundamental query problem. Let K denote a bounded convex polytope in $\mathbb{R}^d$, presented as the intersection of n halfspaces. The objective is to preprocess K so that, given any query point $q \in \mathbb{R}^d$, it is possible to determine efficiently whether q lies in K. Throughout, we assume that d is a fixed constant and K is full-dimensional.

Polytope membership is equivalent in the dual setting to answering halfspace emptiness queries for a set of n points in $\mathbb{R}^d$. In dimensions higher than three, the fastest exact data structure with near-linear space has a query time of roughly $O\big(n^{1-1/\lfloor d/2 \rfloor}\big)$ [29], which is unacceptably high for many applications. Hence, we consider an approximate setting.

Let ε be a positive real parameter, and let $\mathrm{diam}(K)$ denote K's diameter. Given a query point $q \in \mathbb{R}^d$, an *ε-approximate polytope membership query* returns a positive result if $q \in K$, a negative result if the distance from q to its closest point in K is greater than $\varepsilon \cdot \mathrm{diam}(K)$, and it may return either result otherwise.

16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018).
Editor: David Eppstein; Article No. 4; pp. 4:1–4:12
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Polytope membership queries, both exact and approximate, arise in many application areas, such as linear programming and ray-shooting queries [15, 28, 30, 33], nearest neighbor searching and the computation of extreme points [1, 16, 18], collision detection [21], and machine learning [14].

Dudley [20] showed that, for any convex body K in $\mathbb{R}^d$, it is possible to construct an ε-approximating polytope P with $O(1/\varepsilon^{(d-1)/2})$ facets. This bound is asymptotically tight, and is achieved when K is a Euclidean ball. This construction implies a (trivial) data structure for approximate polytope membership problem with space and query time $O(1/\varepsilon^{(d-1)/2})$. It follows from the work of Bentley *et al.* [11] that there is a simple grid-based solution, that answers queries in constant time using space $O(1/\varepsilon^{d-1})$. Arya *et al.* [2,3] present algorithms that achieve a tradeoff between these two extremes, but their data structure provides no improvement over the storage in [11] when the query time is polylogarithmic.

A space-optimal solution for the case of polylogarithmic query time was presented in [7]. It achieves query time $O(\log \frac{1}{\varepsilon})$ with storage $O(1/\varepsilon^{(d-1)/2})$. This paper achieves its efficiency by abandoning the grid- and quadtree-based approaches in favor of an approach based on ellipsoids and a classical structure from convexity theory called a *Macbeath region* [27].

The approach presented in [7] is based on constructing a collection of nested eroded bodies within K and covering the boundaries of these eroded bodies with ellipsoids that are based on Macbeath regions. Queries are answered by shooting rays from a central point in the polytope towards the boundary of K, and tracking an ellipsoid at each level that is intersected by the ray. While it is asymptotically optimal, the data structure and its analysis are complicated by various elements that are artifacts of this ray shooting approach.

In this paper, we present a simpler and more intuitive approach with the same asymptotic complexity as the one in [7]. The key idea is to place the Macbeath regions based on *Delone sets*. A Delone set is a concept from the study of metric spaces. It consists of a set of points that have nice packing and covering properties with respect to the metric balls. Our main result is that any maximal set of disjoint shrunken Macbeath regions defines a Delone set with respect to the Hilbert metric induced on a suitable expansion of the convex body. This observation leads to a simple DAG structure for membership queries. The DAG structure arises from a hierarchy of Delone sets obtained by layering a sequence of expansions of the body. Our results uncover a natural connection between the classical concepts of Delone sets from the theory of metric spaces and Macbeath regions and the Hilbert geometry from the theory of convexity.

2 Preliminaries

In this section we present a number of basic definitions and results, which will be used throughout the paper. We consider the real d-dimensional space, $\mathbb{R}^d$, where d is a fixed constant. Let O denote the origin of $\mathbb{R}^d$. Given a vector $v \in \mathbb{R}^d$, let $\|v\|$ denote its Euclidean length, and let $\langle \cdot, \cdot \rangle$ denote the standard inner product. Given two points $p, q \in \mathbb{R}^d$, the Euclidean distance between them is $\|p - q\|$. For $q \in \mathbb{R}^d$ and $r > 0$, let $B(q, r)$ denote the Euclidean ball of radius r centered at q, and let $B(r) = B(O, r)$.

Let K be a convex body in $\mathbb{R}^d$, represented as the intersection of m closed halfspaces $H_i = \{x \in \mathbb{R}^d : \langle x, v_i \rangle \leq a_i\}$, where a_i is a nonnegative real and $v_i \in \mathbb{R}^d$. The bounding hyperplane for H_i is orthogonal to v_i and lies at distance $a_i/\|v_i\|$ from the origin. The boundary of K will be denoted by ∂K. For $0 < \kappa \leq 1$, we say that K is in *κ-canonical form* if $B(\kappa/2) \subseteq K \subseteq B(1/2)$. Clearly, such a body has a diameter between κ and 1.

It is well known that in $O(m)$ time it is possible to compute a non-singular affine transformation T such that $T(K)$ is in $(1/d)$-canonical form [6, 23]. Further, if a convex body P is within Hausdorff distance ε of $T(K)$, then $T^{-1}(P)$ is within Hausdorff distance at most $d\varepsilon$ of K. (Indeed, this transformation is useful, since the resulting approximation is directionally sensitive, being more accurate along directions where K is skinnier.) Therefore, for the sake of approximation with respect to Hausdorff distance, we may assume that K has been mapped to canonical form, and ε is scaled by a factor of $1/d$. Because we assume that d is a constant, this transformation will only affect the constant factors in our analysis.

A number of our constructions involve perturbing the body K by means of *expansion*, but the exact nature of the expansion is flexible in the following sense. Given $\delta > 0$, let K_δ denote any convex body containing K such that the Hausdorff distance between ∂K and ∂K_δ is $\Theta(\delta \cdot \mathrm{diam}(K))$. For example, if K is in canonical form, K_δ could result as the Minkowski sum of K with another convex body of diameter δ or from a uniform scaling about the origin by δ. Because reducing the approximation parameter by a constant factor affects only the constant factors in our complexity bounds, the use of an appropriate K_δ instead of closely related notions of approximation, like the two just mentioned, will not affect our asymptotic bounds. Given $\delta > 0$, we perturb each H_i to obtain

$$H_{i,\delta} \;=\; \{x \in \mathbb{R}^d : \langle x, \vec{v}_i \rangle \leq a_i + \delta)\}.$$

The associated bounding hyperplane is parallel to that of H_i and translated away from the origin by a distance of $\delta/\|v_i\|$. With that, we define K_δ as the convex polytope $\bigcap_{i=1}^{n} H_{i,\delta}$. To ensure the required bound on the Hausdorff error, we require that $c_1\delta \leq \|v_i\| \leq c_2$ for all i, where c_1 and c_2 are nonnegative reals. The following argument shows that this condition suffices. If $c_1\delta \leq \|v_i\| \leq c_2$, then each bounding halfspace of K is translated away from the origin by a distance of $\delta/\|v_i\| \geq \delta/c_2$, which establishes the lower bound on the Hausdorff distance. Also, each bounding halfspace is translated by a distance of $\delta/\|v_i\| \leq 1/c_1$. Since K, being in canonical form, is nested between balls of radius $\kappa/2$ and $1/2$, this translation of the halfspace is equivalent to a scaling about the origin by a factor of at most $2/c_1\kappa$, which maps each point of K away from the origin by a distance of at most $(2/c_1\kappa)/2 = 1/c_1\kappa$. This establishes the upper bound on the Hausdorff distance.

2.1 Macbeath regions

Our algorithms and data structures will involve packings and coverings by ellipsoids, which will possess the essential properties of Delone sets. These ellipsoids are based on a classical concept from convexity theory, called *Macbeath regions*, which were described first by A. M. Macbeath in a paper on the existence of certain lattice points in a convex body [27]. They have found uses in diverse areas (see, e.g., Bárány's survey [9]).

Given a convex body K, a point $x \in K$, and a real parameter $\lambda \geq 0$, the λ-*scaled Macbeath region* at x, denoted $M_K^\lambda(x)$, is defined to be

$$x + \lambda((K - x) \cap (x - K)).$$

When $\lambda = 1$, it is easy to verify that $M_K^1(x)$ is the intersection of K and the reflection of K around x (see Fig. 1a), and hence it is centrally symmetric about x. $M_K^\lambda(x)$ is a scaled copy of $M_K^1(x)$ by the factor λ about x. We refer to x and λ as the *center* and *scaling factor* of $M_K^\lambda(x)$, respectively. To simplify the notation, when K is clear from the context, we often omit explicit reference in the subscript and use $M^\lambda(x)$ in place of $M_K^\lambda(x)$. When $\lambda < 1$, we say $M^\lambda(x)$ is *shrunken*. When $\lambda = 1$, $M^1(x)$ is *unscaled* and we drop the superscript. Recall that if C^λ is a uniform λ-factor scaling of any bounded, full-dimensional set $C \subset \mathbb{R}^d$, then $\mathrm{vol}(C^\lambda) = \lambda^d \cdot \mathrm{vol}(C)$.

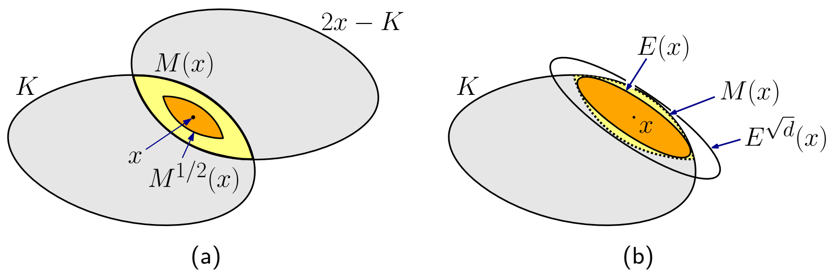

Figure 1 (a) Macbeath regions and (b) Macbeath ellipsoids.

An important property of Macbeath regions, which we call *expansion-containment*, is that if two shrunken Macbeath regions overlap, then an appropriate expansion of one contains the other (see Fig. 2a). The following is a generalization of results of Ewald, Rogers and Larman [22] and Brönnimann, Chazelle, and Pach [13]. Our generalization allows the shrinking factor λ to be adjusted, and shows how to adjust the expansion factor β of the first body to cover an α-scaling of the second body, e.g., the center point only (see Fig. 2b).

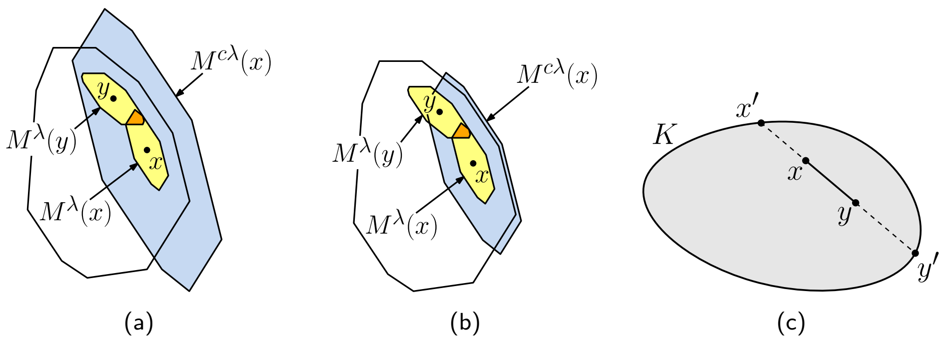

Figure 2 (a)-(b) Expansion-containment per Lemma 1. (c) The Hilbert metric.

▶ **Lemma 1.** *Let $K \subset \mathbb{R}^d$ be a convex body and let $0 < \lambda < 1$. If $x, y \in K$ such that $M^\lambda(x) \cap M^\lambda(y) \neq \emptyset$, then for any $\alpha \geq 0$ and $\beta = \frac{2+\alpha(1+\lambda)}{1-\lambda}$, $M^{\alpha\lambda}(y) \subseteq M^{\beta\lambda}(x)$ (see Fig. 2).*

2.2 Delone sets and the Hilbert metric

An important concept in the context of metric spaces involves coverings and packings by metric balls [19]. Given a metric f over $\mathbb{X}$, a point $x \in \mathbb{X}$, and real $r > 0$, define the ball $B_f(x, r) = \{y \in \mathbb{X} : f(x, y) \leq r\}$. For $\varepsilon, \varepsilon_p, \varepsilon_c > 0$, a set $X \subseteq \mathbb{X}$ is an:

ε-packing: If the balls of radius $\varepsilon/2$ centered at every point of X do not intersect.

ε-covering: If every point of $\mathbb{X}$ is within distance ε of some point of X.

$(\varepsilon_p, \varepsilon_c)$-Delone Set: If X is an ε_p-packing and an ε_c-covering.

Delone sets have been used in the design of data structures for answering geometric proximity queries in metric spaces through the use of hierarchies of nets, such as navigating nets [26], net trees [24], and cover trees [12].

In order to view a collection of Macbeath regions as a Delone set, it will be useful to introduce an underlying metric. The Hilbert metric [25] was introduced over a century ago by David Hilbert as a generalization of the Cayley-Klein model of hyperbolic geometry. A *Hilbert geometry* (K, f_K) consists of a convex domain K in $\mathbb{R}^d$ with the Hilbert distance f_K.

For any pair of distinct points $x, y \in K$, the line passing through them meets ∂K at two points x' and y'. We label these points so that they appear in the order $\langle x', x, y, y' \rangle$ along this line (see Fig. 2c). The Hilbert distance f_K is defined as

$$f_K(x, y) \;=\; \frac{1}{2} \ln \left(\frac{\|x' - y\|}{\|x' - x\|} \frac{\|x - y'\|}{\|y - y'\|} \right).$$

When K is not bounded and either x' or y' is at infinity, the corresponding ratio is taken to be 1. To get some intuition, observe that if x is fixed and y moves along a ray starting at x towards ∂K, $f_K(x, y)$ varies from 0 to ∞.

Hilbert geometries have a number of interesting properties; see the survey by Papadopoulos and Troyanov [32] and the multimedia contribution by Nielsen and Shao [31]. First, f_K can be shown to be a metric. Second, it is invariant under projective transformations.[1] Finally, when K is a unit ball in $\mathbb{R}^d$, the Hilbert distance is equal (up to a constant factor) to the distance between points in the Cayley-Klein model of hyperbolic geometry.

Given a point $x \in K$ and $r > 0$, let $B_H(x, r)$ denote the ball of radius r about x in the Hilbert metric. The following lemma shows that a shrunken Macbeath region is nested between two Hilbert balls whose radii differ by a constant factor (depending on the scaling factor). Thus, up to constant factors in scaling, Macbeath regions and their associated ellipsoids can act as proxies to metric balls in Hilbert space. This nesting was observed by Vernicos and Walsh [34] (for the conventional case of $\lambda = 1/5$), and we present the straightforward generalization to other scale factors. For example, with $\lambda = 1/5$, we have $B_H(x, 0.09) \subseteq M^{1/5}(x) \subseteq B_H(x, 0.21)$ for all $x \in K$.

▶ **Lemma 2.** *Given a convex body $K \subset \mathbb{R}^d$, for all $x \in K$ and any $0 \leq \lambda < 1$,*

$$B_H(x, \frac{1}{2} \ln (1 + \lambda)) \;\subseteq\; M^\lambda(x) \;\subseteq\; B_H\left(x, \frac{1}{2} \ln \frac{1 + \lambda}{1 - \lambda}\right).$$

3 Macbeath regions as Delone sets

Lemma 2 justifies using Macbeath regions as Delone sets. Given a point $x \in K$ and $\delta > 0$, define $M_\delta(x)$ to be the (unscaled) Macbeath region with respect to K_δ, that is, $M_\delta(x) = M_{K_\delta}(x)$. Towards our goal of using Delone sets for approximating convex bodies, we study the behavior of overlapping Macbeath regions at different scales of approximation and establish a bound on the size of such Delone sets. In particular, we consider maximal sets of disjoint shrunken Macbeath regions $M_\delta^\lambda(x)$ defined with respect to K_δ, such that the centers x lie within K; let X_δ denote such a set of centers. The two scale factors used to define the Delone set will be denoted by (λ_p, λ_c), where we assume $0 < \lambda_p < \lambda_c < 1$ are constants. Define $M'_\delta(x) = M_\delta^{\lambda_c}(x)$ and $M''_\delta(x) = M_\delta^{\lambda_p}(x)$.

3.1 Varying the scale

A crucial property of metric balls is how they adapt to changing the resolution at which the domain in question is being modeled. We show that Macbeath regions enjoy a similar property.

[1] This follows from the fact that the argument to the logarithm function is the *cross ratio* of the points (x', x, y, y'), and it is well known that cross ratios are preserved under projective transformations.

▶ **Lemma 3.** *Given a convex body $K \subset \mathbb{R}^d$ and $\lambda, \delta, \varepsilon \geq 0$, for all $x \in K$,*

$$M^{\lambda}_{K_\delta}(x) \;\subseteq\; M^{\lambda}_{K_{(1+\varepsilon)\delta}}(x) \;\subseteq\; M^{(1+\varepsilon)\lambda}_{K_\delta}(x).$$

Proof: The first inclusion is a simple consequence of the fact that enlarging the body can only enlarge the Macbeath regions. To see the second inclusion, it will simplify the notation to translate space by $-x$ so that x now coincides with the origin. Thus, $M_K(x) = K \cap -K$. Recalling our representation from Section 2, we can express K as the intersection of a set of halfspaces $H_i = \{y : \langle y, v_i\rangle \leq a_i\}$. (The translation affects the value of a_i, but not the approximation, because $x \in K$, $a_i \geq 0$.) We can express $M_K(x)$ as the intersection of a set of slabs $\Sigma_i = H_i \cap -H_i$, where each slab is centered about the origin. $M_{K_\delta}(x)$ can be similarly expressed as the intersection of slabs $\Sigma_{i,\delta} = H_{i,\delta} \cap -H_{i,\delta}$, where the defining inequality is $\langle y, v_i\rangle \leq a_i + \delta$. This applies analogously to $M_{K_{(1+\varepsilon)\delta}}(x)$, where the defining inequality is $\langle y, v_i\rangle \leq a_i + (1+\varepsilon)\delta$. Since $a_i \geq 0$, we have $a_i + (1+\varepsilon)\delta \leq (1+\varepsilon)(a_i+\delta)$, which implies that $\Sigma_{i,(1+\varepsilon)\delta} \subseteq (1+\varepsilon)\Sigma_{i,\delta}$. Thus, we have

$$M_{K_{(1+\varepsilon)\delta}}(x) \;=\; \bigcap_{i=1}^{m} \Sigma_{i,(1+\varepsilon)\delta} \;\subseteq\; \bigcap_{i=1}^{m} (1+\varepsilon)\Sigma_{i,\delta} \;=\; M^{(1+\varepsilon)}_{K_\delta}(x).$$

The lemma now follows by applying a scaling factor of λ to both sides. ◻

As we refine the approximation by using smaller values of δ, it is important to bound the number of Macbeath regions at higher resolution that overlap any given Macbeath region at a lower resolution. Our bound is based on a simple packing argument. We will show that the shrunken Macbeath regions $M''_\delta(y)$ that overlap a fixed shrunken Macbeath region at a coarser level of approximation $M'_{s\delta}(x)$, with $s \geq 1$, lie within a suitable constant-factor expansion of $M'_{s\delta}(x)$. Let $Y_{\delta,s}(x)$ denote the set of points y such that $M''_\delta(y)$ are pairwise disjoint and overlap $M'_{s\delta}(x)$. Since these shrunken Macbeath regions are pairwise disjoint, we can bound their number by bounding the ratio of volumes of $M'_{s\delta}(x)$ and $M''_\delta(y)$.

As an immediate corollary of the second inclusion of Lemma 3 we have $\mathrm{vol}(M^{\lambda}_{\delta}(x)) \geq \mathrm{vol}(M^{\lambda}_{s\delta}(x))/s^d$. This allows us to establish an upper bound on the growth rate in the number of Macbeath regions when refining to smaller scales.

▶ **Lemma 4.** *Given a convex body $K \subset \mathbb{R}^d$ and $x \in K$. Then, for constants $\delta \geq 0$, $s \geq 1$ and $Y_{\delta,s}(x)$ as defined above, $|Y_{\delta,s}(x)| = O(1)$.*

Proof: By the first inclusion of Lemma 3, $M'_\delta(y) \subseteq M'_{s\delta}(y)$, and we have $M'_{s\delta}(x) \cap M'_{s\delta}(y) \neq \emptyset$. Next, by applying Lemma 1 (with the roles of x and y swapped) we obtain $M'_{s\delta}(x) = M^{\lambda_c}_{s\delta}(x) \subseteq M^{\beta\lambda_c}_{s\delta}(y)$, with $\alpha = 1$ and $\beta = (3+\lambda_c)/(1-\lambda_c)$.

By definition of X_δ the shrunken Macbeath regions $M''_\delta(y)$ are pairwise disjoint, and so it suffices to bound their volumes with respect to that of $M'_{s\delta}(x)$ to obtain a bound on $|Y_{\delta,s}(x)|$. Applying the corollary to Lemma 3 and scaling, we obtain

$$\mathrm{vol}(M''_\delta(y)) \;\geq\; \frac{1}{s^d}\mathrm{vol}(M''_{s\delta}(y)) \;=\; \left(\frac{\lambda_p}{\beta\lambda_c s}\right)^d \mathrm{vol}(M^{\beta\lambda_c}_{s\delta}(y)) \;\geq\; \left(\frac{\lambda_p}{\beta\lambda_c s}\right)^d \mathrm{vol}(M'_{s\delta}(x)).$$

Thus, by a packing argument the number of children is at most $\left(\frac{\beta\lambda_c s}{\lambda_p}\right)^d = O(1)$. ◻

3.2 Size bound

We bound the cardinality of a maximal set of disjoint shrunken Macbeath regions $M_\delta^\lambda(x)$ defined with respect to K_δ, such that the centers x lie within K; let X_δ denote such a set of centers. This is facilitated by associating each center x with a *cap* of K, where a cap C is defined as the nonempty intersection of the convex body K with a halfspace (see Fig. 3a). Letting h denote the hyperplane bounding this halfspace, the *base* of C is defined as $h \cap K$. The *apex* of C is any point in the cap such that the supporting hyperplane of K at this point is parallel to h. The *width* of C is the distance between h and this supporting hyperplane. Of particular interest is a cap of minimum volume that contains x, which may not be unique. A simple variational argument shows that x is the centroid of the base of this cap [22].

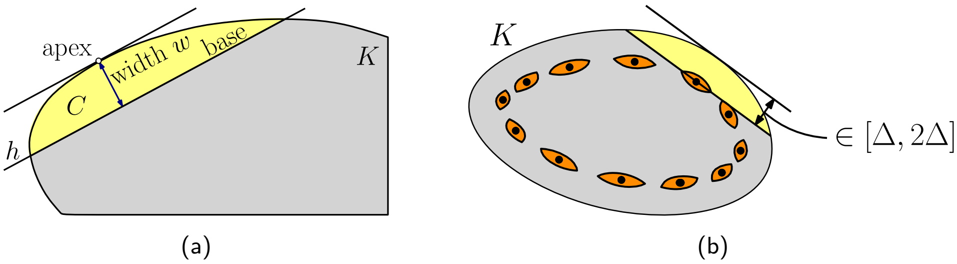

Figure 3 (a) Cap concepts and (b) the economical cap cover.

As each Macbeath region is associated with a cap, we can obtain the desired bound by bounding the number of associated caps. We achieve this by appealing to the so-called *economical cap covers* [10]. The following lemma is a straightforward adaptation of the width-based economical cap cover per Lemma 3.2 of [6].

▶ **Lemma 5.** *Let $K \subset \mathbb{R}^d$ be a convex body in κ-canonical form. Let $0 < \lambda \leq 1/5$ be any fixed constant, and let $\Delta \leq \kappa/12$ be a real parameter. Let $\mathcal{C}$ be a set of caps, whose widths lie between Δ and 2Δ, such that the Macbeath regions $M_K^\lambda(x)$ centered at the centroids x of the bases of these caps are disjoint. Then $|\mathcal{C}| = O(1/\Delta^{(d-1)/2})$ (see Fig. 3a(b)).*

This leads to the following bound on the number of points in X_δ.

▶ **Lemma 6.** *Let $K \subset \mathbb{R}^d$ be a convex body in κ-canonical form, and let X_δ as defined above for some $\delta > 0$ and $0 < \lambda \leq 1/5$. Then, $|X_\delta| = O(1/\delta^{(d-1)/2})$.*

Proof: In order to apply Lemma 5 we will partition the points of X_δ according to the widths of their minimum-volume caps. For $i \geq 0$, define $\Delta_i = c_2 2^i \delta_i$, where c_2 depends on the nature of the the expansion process that yields K_δ. Define $X_{\delta,i}$ to be the subset of points $x \in X_\delta$ such that width of x's minimum cap with respect to K_δ lies within $[\Delta_i, 2\Delta_i]$. By choosing c_2 properly, the Hausdorff distance between K and K_δ is at least $c_2\delta = \Delta_0$, and therefore any cap whose base passes through a point of X_δ has width at least Δ_0. This implies that every point of X_δ lies in some subset $X_{\delta,i}$ for $i \geq 0$.

If a convex body is in κ-canonical form, it follows from a simple geometric argument that for any point x in this body whose minimal cap is of width at least Δ, the body contains a ball of radius $c\Delta$ centered at x, for some constant c (depending on κ and d). If $\Delta_i > \kappa/12$, then $B(x, c\kappa/12) \subseteq K_\delta$ for all $x \in X_{\delta,i}$. It follows that $B(x, c\kappa/12) \subseteq M_\delta(x)$ implying that $\mathrm{vol}(M_\delta^\lambda(x)) \geq \lambda^d \cdot \mathrm{vol}(B(c\kappa/12))$ which is $\Omega(1)$ as c, κ and λ are all constants. By a simple packing argument $|X_{i,j}| = O(1)$. There are at most a constant number of levels for which $\Delta_j > \kappa/12$, and so the overall contribution of these subsets is $O(1)$.

Henceforth, we may assume that $\Delta_j \leq \kappa/12$. Since $\lambda \leq 1/5$, we apply Lemma 5 to obtain the bound $|X_{\delta,i}| = O(1/\Delta_i^{(d-1)/2})$. (There is a minor technicality here. If δ becomes sufficiently large, K_δ may not be in κ-canonical form because its diameter is too large. Because $\delta = O(1)$ and hence $\mathrm{diam}(K_\delta) = O(1)$, we may scale it back into canonical form at the expense of increasing the constant factors hidden in the asymptotic bound.) Thus, up to constant factors, we have

$$|X_\delta| \;=\; \sum_{i\geq 0} |X_{\delta,i}| \;=\; \sum_{i\geq 0} O\left(\frac{1}{\Delta_i}\right)^{\frac{d-1}{2}} \;=\; \sum_{i\geq 0} O\left(\frac{1}{c_2 2^i \delta}\right)^{\frac{d-1}{2}} \;=\; O\left(\left(\frac{1}{\delta}\right)^{\frac{d-1}{2}}\right).$$

◻

4 Macbeath ellipsoids

For the sake of efficient computation, it will be useful to approximate Macbeath regions by shapes of constant combinatorial complexity. We have opted to use ellipsoids. (Note that bounding boxes [1] could be used instead, and may be preferred in contexts where polytopes are preferred.)

Given a Macbeath region, define its associated *Macbeath ellipsoid* $E_K^\lambda(x)$ to be the maximum-volume ellipsoid contained within $M_K^\lambda(x)$ (see Fig. 1b). Clearly, this ellipsoid is centered at x and $E_K^\lambda(x)$ is an λ-factor scaling of $E_K^1(x)$ about x. It is well known that the maximum-volume ellipsoid contained within a convex body is unique, and Chazelle and Matoušek showed that it can be computed for a convex polytope in time linear in the number of its bounding halfspaces [17]. By John's Theorem (applied in the context of centrally symmetric bodies) it follows that $E_K^\lambda(x) \subseteq M_K^\lambda(x) \subseteq E_K^{\lambda\sqrt{d}}(x)$ [8].

Given a point $x \in K$ and $\delta > 0$, define $M_\delta(x)$ to be the (unscaled) Macbeath region with respect to K_δ (as defined in Section 2), that is, $M_\delta(x) = M_{K_\delta}(x)$. Let $E_\delta(x)$ denote the maximum volume ellipsoid contained within $M_\delta(x)$. As $M_\delta(x)$ is symmetric about x, $E_\delta(x)$ is centered at x. For any $\lambda > 0$, define $M_\delta^\lambda(x)$ and $E_\delta^\lambda(x)$ to be the uniform scalings of $M_\delta(x)$ and $E_\delta(x)$, respectively, about x by a factor of λ. By John's Theorem, we have

$$E_\delta^\lambda(x) \;\subseteq\; M_\delta^\lambda(x) \;\subseteq\; E_\delta^{\lambda\sqrt{d}}(x). \tag{1}$$

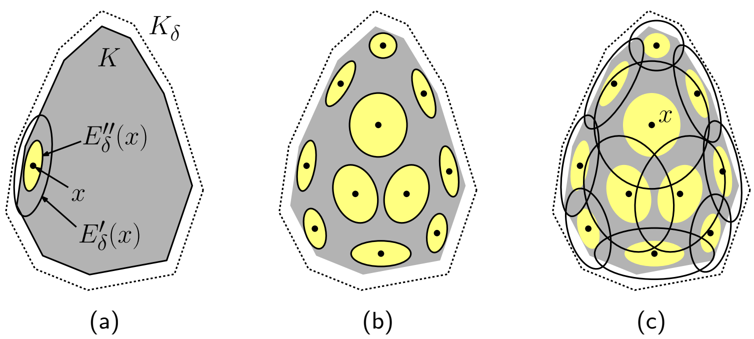

Figure 4 A Delone set for a convex body. (Not drawn to scale.)

Two particular scale factors will be of interest to us. Define $M'_\delta(x) = M_\delta^{1/2}(x)$ and $M''_\delta(x) = M_\delta^{\lambda_0}(x)$, where $\lambda_0 = 1/(4\sqrt{d}+1)$. Similarly, define $E'_\delta(x) = E_\delta^{1/2}(x)$ and $E''_\delta(x) = E_\delta^{\lambda_0}(x)$ (see Fig. 4(a)). Given a fixed δ, let X_δ be any maximal set of points, all lying within K, such that the ellipsoids $E''_\delta(x)$ are pairwise disjoint for all $x \in X_\delta$.

These ellipsoids form a packing of K_δ (see Fig. 4(b)). The following lemma shows that their suitable expansions cover K while being contained within K_δ (see Fig. 4(c)).

▶ **Lemma 7.** *Given a convex body K in $\mathbb{R}^d$ and a set X_δ as defined above for $\delta > 0$,*

$$K \subseteq \bigcup_{x \in X_\delta} E'_\delta(x) \subseteq K_\delta.$$

Proof: To establish the first inclusion, consider any point $y \in K$. Because X_δ is maximal, there exists $x \in X_\delta$ such that $E''_\delta(x) \cap E''_\delta(y)$ is nonempty. By containment, $M''_\delta(x) \cap M''_\delta(y)$ is also nonempty. By Lemma 1 (with $\alpha = 0$), it follows that $y \in M^\lambda_\delta(x)$, where

$$\lambda = \frac{2\lambda_0}{1-\lambda_0} = \frac{2/(4\sqrt{d}+1)}{1-1/(4\sqrt{d}+1)} = \frac{2}{4\sqrt{d}} = \frac{1}{2\sqrt{d}}.$$

By applying Eq. (1) (with $\lambda = 1/(2\sqrt{d})$), we have $M^{1/(2\sqrt{d})}_\delta(x) \subseteq E^{1/2}_\delta(x) = E'_\delta(x)$, and therefore $y \in E'_\delta(x)$. Thus, we have shown that an arbitrary point $y \in K$ is contained in the ellipsoid $E'_\delta(x)$ for some $x \in X_\delta$, implying that the union of these ellipsoids covers K. The second inclusion follows from $E'_\delta(x) \subseteq M'_\delta(x) \subseteq M_\delta(x) \subseteq K_\delta$ for any $x \in X_\delta \subseteq K$. ◻

In conclusion, if we treat the scaling factor λ in $E^\lambda(x)$ as a proxy for the radius of a metric ball, we have shown that X_δ is a $(2\lambda_0, 1/2)$-Delone set for K. By Lemma 2 this is also true in the Hilbert metric over K_δ up to a constant factor adjustment in the radii. (Note that the scale of the Hilbert balls does not vary with δ. What varies is the choice of the expanded body K_δ defining the metric.)

By John's Theorem, Macbeath regions and Macbeath ellipsoids differ by a constant scaling factor, both with respect to enclosure and containment. We remark that all the results of the previous two sections hold equally for Macbeath ellipsoids. We omit the straightforward, but tedious, details.

▶ Remark. All results from previous section on scaled Macbeath regions apply to scaled Macbeath ellipsoids subject to appropriate modifications of the constant factors.

5 Approximate polytope membership (APM)

The Macbeath-based Delone sets developed above yield a simple data structure for answering ε-APM queries for a convex body K. We assume that K is represented as the intersection of m halfspaces. We may assume that in $O(m)$ time it has been transformed into κ-canonical form, for $\kappa = 1/d$. Throughout, we will assume that Delone sets are based on the Macbeath ellipsoids $E''_\delta(x)$ for packing and $E'_\delta(x)$ for coverage (defined in Section 4).

Our data structure is based on a hierarchy of Delone sets of exponentially increasing accuracy. Define $\delta_0 = \varepsilon$, and for any integer $i \geq 0$, define $\delta_i = 2^i \delta_0$. Let X_i denote a Delone set for K_{δ_i}. By Lemma 7, we may take X_i to be any maximal set of points within K such that the packing ellipsoids $E''_\delta(x)$ are pairwise disjoint. Let $\ell = \ell_\varepsilon$ be the smallest integer such that $|X_\ell| = 1$. We will show below that $\ell = O(\log 1/\varepsilon)$.

Given the sets $\langle X_0, \ldots, X_\ell \rangle$, we build a rooted, layered DAG structure as follows. The nodes of level i correspond 1–1 with the points of X_i. The leaves reside at level 0 and the root at level ℓ. Each node $x \in X_i$ is associated with two things. The first is its *cell*, denoted $\mathrm{cell}(x)$, which is the covering ellipsoid $E'_\delta(x)$ (the larger hollow ellipsoids shown in Fig. 5). The second, if $i > 0$, is a set of *children*, denoted $\mathrm{ch}(x)$, which consists of the points $y \in X_{i-1}$ such that $\mathrm{cell}(x) \cap \mathrm{cell}(y) \neq \emptyset$.

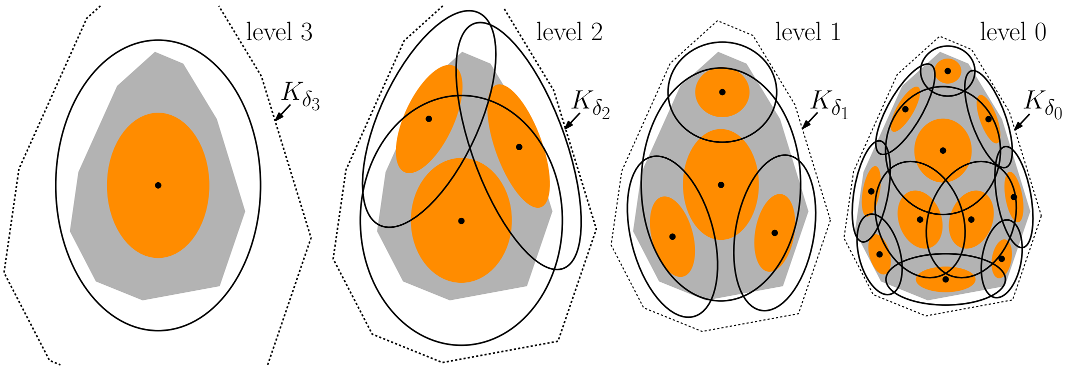

Figure 5 Hierarchy of ellipsoids for answering APM queries.

To answer a query q, we start at the root and iteratively visit any one node $x \in X_i$ at each level of the DAG, such that $q \in \mathrm{cell}(x)$. We know that if q lies within K, such an x must exist by the covering properties of Delone sets, and further at least one of x's children contains q. If q does not lie within any of the children of the current node, the query algorithm terminates and reports (without error) that $q \notin K$. Otherwise the search eventually reaches a node $x \in X_0$ at the leaf level whose cell contains q. Since $\mathrm{cell}(x) \subseteq K_{\delta_0} = K_\varepsilon$, this cell serves as a witness to q's approximate membership within K.

In order to bound the space and query time, we need to bound the total space used by the data structure and the time to process each node in the search, which is proportional to the number of its children. Building upon Lemmas 4 and 6, we have our main result.

▶ **Theorem 8.** *Given a convex body K and $\varepsilon > 0$, there exists a data structure of space $O(1/\varepsilon^{(d-1)/2})$ that answers ε-approximate polytope membership queries in time $O(\log 1/\varepsilon)$.*

Since the expansion factors δ_i grow exponentially from ε to a suitably large constant, it follows that the height of the tree is logarithmic in $1/\varepsilon$, which is made formal below.

▶ **Lemma 9.** *The DAG structure described above has height $O(\log 1/\varepsilon)$.*

Proof: Let c_2 be an appropriate constant, and let $\ell = \lceil \log_2(2/c_2\varepsilon)\rceil = O(\log 1/\varepsilon)$. Depending the nature of the expanded body K_δ, the constant c_2 can be chosen so the Hausdorff distance between K and K_{δ_ℓ} is at least $c_2\delta_\ell = c_2 2^\ell \varepsilon \geq 2$. Because K is in κ-canonical form, it is contained within a unit ball centered at the origin. Therefore, K_{δ_ℓ} contains a ball of radius two centered at the origin, which implies that the Macbeath ellipsoid $E'_{\delta_\ell}(O)$ (which is scaled by $1/2$) contains the unit ball and so contains K. Thus, (assuming that the origin is added first to the Delone set) level ℓ of the DAG contains a single node. ◻

By Lemma 4, each node has $O(1)$ children and $\delta_i = 2^i\delta_0 = 2^i\varepsilon$, we obtain the following space bound by summing $|X_i|$ for $0 \leq i \leq \ell$.

▶ **Lemma 10.** *The storage required by the DAG structure described above is $O(1/\varepsilon^{(d-1)/2})$.*

As mentioned above, by combining Lemmas 4 with 6, it follows that the query time is $O(\log 1/\varepsilon)$ and by Lemma 10 the total space is $O(1/\varepsilon^{(d-1)/2})$, which establish Theorem 8.

While our focus has been on demonstrating the existence of a simple data structure derived from Delone sets, we note that it can be constructed by well-established techniques. While obtaining the best dependencies on ε in the construction time will likely involve fairly sophisticated methods, as seen in the paper of Arya et al. [5], the following shows that there is a straightforward construction.

▶ **Lemma 11.** *Given a convex body $K \subset \mathbb{R}^d$ represented as the intersection of m halfspaces and $\varepsilon > 0$, the above DAG structure for answering ε-APM queries can be computed in time $O(m + 1/\varepsilon^{O(d)})$, where the constant in the exponent does not depend on ε or d.*

Proof: First, we transform K into canonical form, and replace it with an $\frac{\varepsilon}{2}$-approximation K' of itself. This can be done in $O(m + 1/\varepsilon^{O(d)})$, so that K' is bounded by $O(1/\varepsilon^{(d-1)/2})$ halfspaces (see, e.g., [4]). We then build the data structure to solve APM queries to an accuracy of $(\varepsilon/2)$, so that the total error is ε.

Because the number of nodes increases exponentially as we descend to the leaf level, the most computationally intensive aspect of the remainder of the construction is computing the set X_0, a maximal subset of K whose packing ellipsoids $E''_{\delta_0}(x)$ are pairwise disjoint. To discretize the construction of X_0, we observe that by our remarks at the start of Section 2, the Hausdorff distance between K and K_{δ_0} is $\Omega(\delta_0) = \Omega(\varepsilon)$. It follows that each of the ellipsoids $E''_{\delta_0}(x)$ contains a ball of radius $\Omega(\lambda_0 \varepsilon) = \Omega(\varepsilon)$. We restrict the points of X_0 to come from the vertices of a square grid whose side length is half this radius. Since K is in canonical form, it suffices to generate $O(1/\varepsilon^{O(d)})$ grid points. By decreasing the value of ε slightly (by a constant factor), it is straightforward to show that any Delone set can be perturbed so that its centers lie on this grid.

Each Macbeath ellipsoid can be computed in time linear in the number of halfspaces bounding K', which is $O(1/\varepsilon^{O(d)})$ [17]. The maximal set is computed by brute force, repeatedly selecting a point x from the grid, computing $E''_{\delta_0}(x)$, and marking the points of the grid that it covers until all points interior to K are covered. The overall running time is dominated by the product of the number of grid points and the $O(1/\varepsilon^{O(d)})$ time to compute each Macbeath ellipsoid. ◻

References

1 P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. Approximating extent measures of points. *J. Assoc. Comput. Mach.*, 51:606–635, 2004.

2 S. Arya, G. D. da Fonseca, and D. M. Mount. Approximate polytope membership queries. In *Proc. 43rd Annu. ACM Sympos. Theory Comput.*, pages 579–586, 2011.

3 S. Arya, G. D. da Fonseca, and D. M. Mount. Polytope approximation and the Mahler volume. In *Proc. 23rd Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 29–42, 2012.

4 S. Arya, G. D. da Fonseca, and D. M. Mount. Approximate polytope membership queries. *SIAM J. Comput.*, 2017. To appear.

5 S. Arya, G. D. da Fonseca, and D. M. Mount. Near-optimal ε-kernel construction and related problems. In *Proc. 33rd Internat. Sympos. Comput. Geom.*, pages 10:1—10:15, 2017.

6 S. Arya, G. D. da Fonseca, and D. M. Mount. On the combinatorial complexity of approximating polytopes. *Discrete & Computational Geometry*, 58(4):849–870, 2017.

7 S. Arya, G. D. da Fonseca, and D. M. Mount. Optimal approximate polytope membership. In *Proc. 28th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 270–288, 2017.

8 K. Ball. An elementary introduction to modern convex geometry. In S. Levy, editor, *Flavors of Geometry*, pages 1–58. Cambridge University Press, 1997. (MSRI Publications, Vol. 31).

9 I. Bárány. The technique of M-regions and cap-coverings: A survey. *Rend. Circ. Mat. Palermo*, 65:21–38, 2000.

10 I. Bárány and D. G. Larman. Convex bodies, economic cap coverings, random polytopes. *Mathematika*, 35:274–291, 1988.

11 J. L. Bentley, M. G. Faust, and F. P. Preparata. Approximation algorithms for convex hulls. *Commun. ACM*, 25(1):64–68, 1982.

12 A. Beygelzimer, S. Kakade, and J. Langford. Cover trees for nearest neighbor. In *Proc. 23rd Internat. Conf. on Machine Learning*, pages 97–104, 2006.

13 H. Brönnimann, B. Chazelle, and J. Pach. How hard is halfspace range searching. *Discrete Comput. Geom.*, 10:143–155, 1993.

14 C. J. C. Burges. A tutorial on support vector machines for pattern recognition. *Data Min. Knowl. Discov.*, 2(2):121–167, 1998.

15 T. M. Chan. Fixed-dimensional linear programming queries made easy. In *Proc. 12th Annu. Sympos. Comput. Geom.*, pages 284–290, 1996.

16 T. M. Chan. Output-sensitive results on convex hulls, extreme points, and related problems. *Discrete Comput. Geom.*, 16:369–387, 1996.

17 B. Chazelle and J. Matoušek. On linear-time deterministic algorithms for optimization problems in fixed dimension. *J. Algorithms*, 21:579–597, 1996.

18 K. L. Clarkson. An algorithm for approximate closest-point queries. In *Proc. Tenth Annu. Sympos. Comput. Geom.*, pages 160–164, 1994.

19 K. L. Clarkson. Building triangulations using ε-nets. In *Proc. 38th Annu. ACM Sympos. Theory Comput.*, pages 326–335, 2006.

20 R. M. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *J. Approx. Theory*, 10(3):227–236, 1974.

21 J. Erickson, L. J. Guibas, J. Stolfi, and L. Zhang. Separation-sensitive collision detection for convex objects. In *Proc. Tenth Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 327–336, 1999.

22 G. Ewald, D. G. Larman, and C. A. Rogers. The directions of the line segments and of the r-dimensional balls on the boundary of a convex body in Euclidean space. *Mathematika*, 17:1–20, 1970.

23 S. Har-Peled. A replacement for Voronoi diagrams of near linear size. In *Proc. 42nd Annu. IEEE Sympos. Found. Comput. Sci.*, pages 94–103, 2001.

24 S. Har-Peled and M. Mendel. Fast construction of nets in low dimensional metrics, and their applications. *SIAM J. Comput.*, 35:1148–1184, 2006.

25 D. Hilbert. Ueber die gerade linie als kürzeste verbindung zweier punkte. *Mathematische Annalen*, 46:91–96, 1895.

26 R. Krauthgamer and J. R. Lee. Navigating nets: Simple algorithms for proximity search. In *Proc. 15th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 798–807, 2004.

27 A. M. Macbeath. A theorem on non-homogeneous lattices. *Ann. of Math.*, 56:269–293, 1952.

28 J. Matoušek and O. Schwarzkopf. On ray shooting in convex polytopes. *Discrete Comput. Geom.*, 10:215–232, 1993.

29 J. Matoušek. Reporting points in halfspaces. *Comput. Geom. Theory Appl.*, 2:169–186, 1992.

30 J. Matoušek. Linear optimization queries. *J. Algorithms*, 14(3):432–448, 1993.

31 F. Nielsen and L. Shao. On Balls in a Hilbert Polygonal Geometry (Multimedia Contribution). In *Proc. 33rd Internat. Sympos. Comput. Geom.*, pages 67:1–67:4, 2017.

32 A. Papadopoulos and M. Troyanov. *Handbook of Hilbert Geometry*. European Mathematical Society, 2014.

33 E. A. Ramos. Linear programming queries revisited. In *Proc. 16th Annu. Sympos. Comput. Geom.*, pages 176–181, 2000.

34 C. Vernicos and C. Walsh. Flag-approximability of convex bodies and volume growth of Hilbert geometries. HAL Archive (hal-01423693i), 2016. URL: https://hal.archives-ouvertes.fr/hal-01423693.