

Polytope Approximation and the Mahler Volume

Sunil Arya*

Department of Computer Science and Engineering
The Hong Kong University of
Science and Technology
Clear Water Bay, Kowloon, Hong Kong
arya@cse.ust.hk

Guilherme D. da Fonseca†

Departamento de Informática Aplicada
Universidade Federal do Estado
do Rio de Janeiro (UNIRIO)
Rio de Janeiro, Brazil
fonseca@uniriotec.br

David M. Mount‡

Department of Computer Science and
Institute for Advanced Computer Studies
University of Maryland
College Park, Maryland 20742
mount@cs.umd.edu

Abstract

The problem of approximating convex bodies by polytopes is an important and well studied problem. Given a convex body K in R^d , the objective is to minimize the number of vertices (alternatively, the number of facets) of an approximating polytope for a given Hausdorff error ε . Results to date have been of two types. The first type assumes that K is smooth, and bounds hold in the limit as ε tends to zero. The second type requires no such assumptions. The latter type includes the well known results of Dudley (1974) and Bronshteyn and Ivanov (1976), which show that in spaces of fixed dimension, $O((\text{diam}(K)/\varepsilon)^{(d-1)/2})$ vertices (alt., facets) suffice. Our results are of this latter type.

In our first result, under the assumption that the width of the body in any direction is at least ε , we strengthen the above bound to $\tilde{O}(\sqrt{\text{area}(K)}/\varepsilon^{(d-1)/2})$. This is never worse than the previous bound (by more than logarithmic factors) and may be significantly better for skinny bodies. Our analysis exploits an interesting analogy with a classical concept from the theory of convexity, called the Mahler volume. This is a dimensionless quantity that involves the product of the volumes of a convex body and its polar dual.

In our second result, we apply the same machinery to improve upon the best known bounds for answering

ε -approximate polytope membership queries. Given a convex polytope P defined as the intersection of halfspaces, such a query determines whether a query point q lies inside or outside P , but may return either answer if q 's distance from P 's boundary is at most ε . We show that, without increasing storage, it is possible to reduce the best known search times for ε -approximate polytope membership significantly. This further implies improvements to the best known search times for approximate nearest neighbor searching in spaces of fixed dimension.

1 Introduction

Approximating convex bodies by polytopes is a fundamental problem, which has been extensively studied in the literature. (See [6] for a recent survey.) The problem is to determine the minimum number of vertices n (alternatively, the minimum number of facets) of an approximating polytope for a given error $\varepsilon > 0$. Error is commonly measured through the Hausdorff distance [6]. Given a convex body K in d -dimensional space, we say that a polytope P ε -approximates K if the Hausdorff distance between K and P is at most ε . Throughout this paper, we will restrict attention to the Hausdorff metric and to Euclidean space of constant dimension d .

Although much work has been done on the optimization problem of finding the “best” approximating polytope for a specific input according to various criteria (see, for example, [9, 13, 17, 22]), our interest is in establishing bounds on the combinatorial complexity of approximating general convex bodies. Approximation

*Research supported by the Research Grants Council of Hong Kong, China under project number 610110.

†Research supported by CNPq and FAPERJ grants.

‡Research supported by NSF grant CCR-0635099 and ONR grant N00014-08-1-1015.

bounds are typically expressed as the number n of vertices or facets of a convex polytope needed to achieve a certain approximation error ε for a given convex body K . Results are of two common types. In both cases, it is shown that there exists $\varepsilon_0 > 0$, and the bounds hold for all $\varepsilon \leq \varepsilon_0$. In the first type, which we call *nonuniform bounds*, the value of ε_0 depends on K (for example, on K 's maximum curvature). Such bounds are often stated as holding “in the limit” as ε approaches zero, or equivalently as n approaches infinity. In the second type, which we call *uniform bounds*, the value of ε_0 is independent of K .

As an example of a nonuniform bound, consider the following result due to Gruber [14]. Let ∂K denote K 's boundary. For a point $x \in \partial K$, let $\kappa(x)$ denote the Gaussian curvature at x , and let $d\sigma(x)$ denote the ordinary differential surface element at x . Under certain smoothness conditions on K , Gruber showed that there exists an approximating polytope with n vertices, where

$$n = O\left(\frac{1}{\varepsilon^{(d-1)/2}} \int_{\partial K} \sqrt{\kappa(x)} d\sigma(x)\right),$$

as $\varepsilon \rightarrow 0$. The constant hidden in the O -notation depends only on dimension. This bound also applies to the number of facets n in an approximating polytope. Recently, Clarkson [10] has extended these results and shown a similar bound on the triangulation complexity for more general manifolds that may not represent the boundary of a convex body. Other examples of nonuniform bounds can be found in [15, 21, 26, 27].

In contrast, examples of uniform bounds include Dudley [12] and Bronshteyn and Ivanov [5]. These bounds hold without any smoothness assumptions, and they hold even for polytopes. Dudley showed that, for $\varepsilon \leq 1$, any convex body K can be ε -approximated by a polytope P with $O((\text{diam}(K)/\varepsilon)^{(d-1)/2})$ facets. Bronshteyn and Ivanov showed the same bound holds for the number of vertices. These bounds are tight up to constant factors (e.g., when K is a ball). Again, constants hidden in the O -notation depend only on d . These results have many applications, for example, in the construction of coresets [1]. In this paper, we consider only uniform bounds.

1.1 Our Results We present two main results, which arise by replacing the approximation bounds of both Dudley and Bronshteyn and Ivanov with bounds that are more sensitive to the local geometry of the convex body. Our first result shows that the existing bounds can be significantly improved for the case when K has directional width smaller than the diameter along some directions, that is, when K is skinny. It seems intuitive that skinny bodies should be easier to approximate,

but we are unaware of any uniform bounds in the literature that capture this fact. Under the reasonable assumption¹ that the width of the body in any direction is at least ε , we show that it is possible to construct an ε -approximating polytope P in which the number of facets² is $\tilde{O}(\sqrt{\text{area}(K)}/\varepsilon^{(d-1)/2})$, where $\text{area}(K)$ denotes the surface area of K . Although we do not include it here, a similar bound can be proved for the number of vertices. Since $\text{area}(K)$ is bounded by $O(\text{diam}(K)^{d-1})$, our bound is as good as Dudley's for fat convex bodies (up to logarithmic factors), and it is superior for skinny ones. These results are presented in Section 4, with the main result given in Theorem 4.1.

Our second result involves approximate polytope membership queries. Given a convex polytope P , defined as the intersection of halfspaces, an ε -approximate polytope membership query determines whether a query point q lies inside or outside P , but it may return either answer if q 's distance from P 's boundary is at most ε . Polytope membership queries find applications in many geometric areas, such as linear programming queries, ray shooting, nearest neighbor searching, and the computation of convex hulls [7, 8, 18, 20, 23]. In dimension $d \leq 3$, it is possible to build a data structure of linear size that can answer such queries in logarithmic time [11]. In higher dimensions, however, all exact data structures with roughly linear space take $\tilde{O}(n^{1-1/d/2})$ query time [19], which, except in small dimensions, is little better than brute-force search.

Assuming that P has unit diameter, a straightforward application of Dudley's method allows such queries to be answered in $O(1/\varepsilon^{(d-1)/2})$ storage and query time (by a brute-force enumeration of the approximating halfspaces). We recently showed that there is a very simple hierarchical search algorithm that achieves the same storage bound but improves the query time to $O(1/\varepsilon^{(d-1)/4})$, thus reducing the exponent by half [2]. In this paper, we show that, with the same storage bounds, it is possible to again reduce the exponent in the query time by half, to $O(1/\varepsilon^{(d-1)/8})$. We also present

¹Without this assumption it is impossible to achieve area-based bounds. For example, given a fat d -dimensional convex body, embed it in dimension $d+2$ and take its Minkowski sum with a ball of radius $\delta \ll \varepsilon$. The surface area of the resulting body can be made arbitrarily small as a function of δ , but clearly the complexity of any approximating body is not smaller than for the original body. Conversely, if the body is of width less than ε in some direction, then by projecting the body onto a hyperplane orthogonal to this direction, we can reduce the problem to a convex approximation problem in one lower dimension. This can be repeated until the body's width is sufficiently large, after which our bounds apply.

²Throughout, we treat ε as an asymptotic quantity and use \tilde{O} to indicate that logarithmic factors in ε are ignored.

a space-time tradeoff that improves query times to the square root of the previous bounds throughout the spectrum. The results are summarized in Figure 1(a).

It is notable that the improvement involves no change to the simple data structure of [2], just a better analysis of its complexity. These results are presented in Section 5, and the main result is given in Theorem 5.1. As shown in [2], these results provide immediate improvements to the best space-time tradeoffs for approximate nearest neighbor searching (see Figure 1(b)).

1.2 Overview of Methods Our results stem from a locally sensitive approach to convex approximation. Our area-based bound on the complexity of convex body approximation comes about through a consideration of the complexity of approximating small patches on the boundary of the convex body, called caps. Given a halfspace H whose bounding hyperplane intersects K , we define a *cap* to be the portion of ∂K that is contained within H . The *width* of the cap is the maximum distance between any point on the cap and the bounding hyperplane of H , and an ε -*cap* is a cap of width ε . A set of points is said to *stab* all ε -caps if every ε -cap has a nonempty intersection with the set. Bronshteyn and Ivanov [5] observed that the convex hull of any set of points that stabs all ε -caps yields an ε -approximation of K . Assume that K has unit diameter, and let S be the $(d - 1)$ -dimensional sphere of radius $\frac{1}{2}$ centered at the origin, which we call the *Dudley hypersphere*. (Any sufficiently large constant radius suffices.) We may associate any ε -cap C with the surface patch on S consisting of points whose closest point on ∂K lies on C ; we call this the *Voronoi patch* of C .

Our key geometric observation is that the product of the area of any ε -cap and its associated Voronoi patch must be large. Intuitively, if the surface area of the ε -cap is small, then its curvature must be high, and so its associated Voronoi patch must have relatively large area. We exploit this observation to efficiently stab all the ε -caps using a two-pronged strategy. We stab ε -caps that have large surface area by selecting points randomly on the boundary of K using density proportional to the surface area. We stab the ε -caps that have small surface area by selecting points randomly on the Dudley hypersphere (again using density proportional to the surface area), and then choosing their nearest neighbor on ∂K . We will see in Section 4 that this allows us to stab all the ε -caps with $\tilde{O}(\sqrt{\text{area}(K)}/\varepsilon^{(d-1)/2})$ points.

Our approach to answering polytope membership queries for a convex body K is based on a quadtree search, where the basic problem involves approximating the intersection of the boundary of K with a single quadtree box. Our improved bounds for polytope mem-

bership queries follow as a consequence of establishing improved bounds for such approximations, whenever the diameter of the quadtree box is at least $\sqrt{\varepsilon}$. Roughly speaking, we show that if A_1 denotes the area of $\partial K \cap b$, and A_2 denotes the area of the associated Voronoi patch, then $\partial K \cap b$ can be ε -approximated using $O(\sqrt{A_1 A_2}/\varepsilon^{(d-1)/2})$ facets. From this it is not hard to show that for a box b of diameter $\sqrt{\varepsilon}$, if Dudley’s method uses m facets to ε -approximate $\partial K \cap b$, then $O(\sqrt{m})$ facets suffice. We then exploit this fact in presenting an improved analysis of the search algorithm of [2].

The geometric underpinnings needed to show that the product of areas is bounded is presented in Section 3. We introduce a dual notion of cap, called a *dual cap*, and we explore the relationship between the areas of surface caps and their Voronoi patches. We demonstrate an interesting parallel between these concepts and a classical notion from the theory of convexity, called the *Mahler volume* [4, 25]. The Mahler volume of a convex body is a dimensionless quantity that involves the product of the body’s volume and the volume of its polar body. We demonstrate that caps and their Voronoi patches exhibit a similar polar relationship.

The improvements presented here do not arise simply by a more careful application of the analysis given in [2]. The analysis given there relied on a direct application of Dudley’s and Bronshteyn and Ivanov’s bounds. The key to the improvements of this paper is the development of uniform bounds that are sensitive to local geometry. Sensitivity to local geometry has been applied in asymptotic contexts (under the assumption of smoothness) by Gruber [14] and Clarkson [10]. It has also been applied for uniform bounds, but only in the two-dimensional setting by Rote [24]. The techniques developed here provide powerful tools for analyzing the complexity of approximating arbitrary convex bodies on a local level, and we believe that these techniques will be of wider applicability.

2 Preliminaries

Let K denote a convex body in \mathbb{R}^d , that is, a nonempty, closed, bounded convex set. Let ∂K denote K ’s boundary. Given $\delta > 0$, we say that a convex body is δ -*smooth* if at each point $q \in \partial K$ there exists a ball of radius δ that lies entirely within K and contains q on its boundary. We say that K is *smooth* if it is δ -smooth for some $\delta > 0$. If K is smooth, then for any $q \in \partial K$, there is a unique supporting hyperplane for K containing q , which we denote by $h(q)$.

Given any bounded set X in \mathbb{R}^d , which for us will be a convex body, we use $\text{vol}(X)$ to denote its d -dimensional Hausdorff measure. Given a $(d - 1)$ -dimensional manifold Y in \mathbb{R}^d , for example, a surface

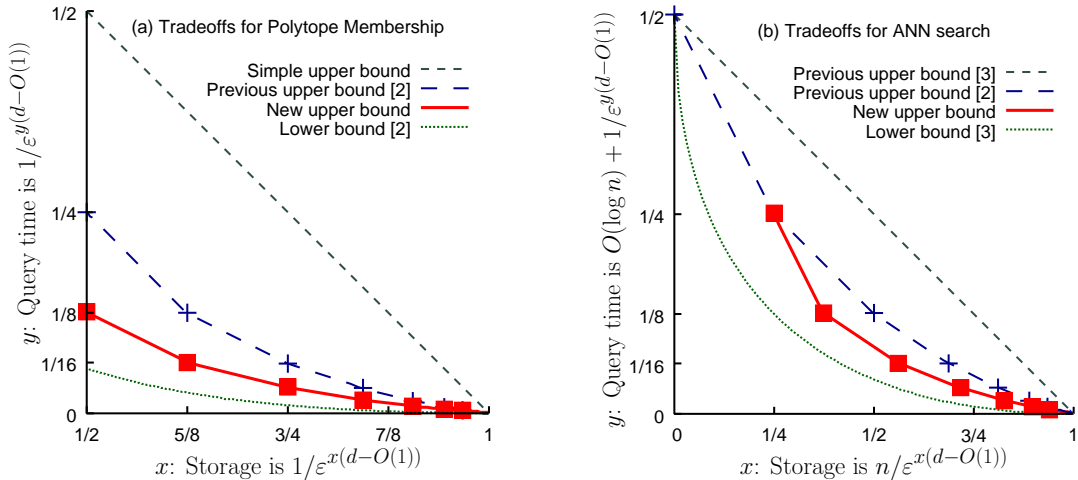


Figure 1: The multiplicative factor in the exponent of the $1/\varepsilon$ term for (a) polytope membership queries and (b) approximate nearest neighbor (ANN) queries. The $O(1)$ term in the exponent corresponds to a constant that does *not* depend on d .

patch on a convex body, we let $\text{area}(Y)$ denote its $(d-1)$ -dimensional Hausdorff measure. Given a convex body X in \mathbb{R}^d , we use $\text{area}(X)$ as a shorthand for $\text{area}(\partial X)$. We say that a polytope P ε -approximates K if the Hausdorff distance between K and P is at most ε .

Let H be a halfspace whose bounding hyperplane intersects K . The intersection $K \cap H$ is called the *cap* generated by H . Since we are interested in the surface of K , henceforth we will use the term *cap* in a nonstandard sense as the portion of the boundary of the cap that coincides with ∂K (in essence, the outer skin of the cap). Let q denote a point of the cap of maximum distance from ∂H . By local maximality, $h(q)$ and ∂H are parallel to each other, and the distance between them is called the *width* of the cap. A cap of width ε is called an ε -cap. We refer to q as the cap's *apex*, and $K \cap \partial H$ as the cap's *base*, which we denote by Γ_q . By smoothness, for each $q \in \partial K$, there is a unique ε -cap having q as its apex, which we denote by C_q (see Fig 2(a)).




Figure 2: Caps and dual caps.

Given any point q external to K , the dual cap³ generated by q is the topological closure of the intersection of all halfspaces that contain K but not q . As before, since we are interested in the surface of K , henceforth we use the term *dual cap*, denoted by D_q , to refer to the portion of the boundary of the dual cap that coincides with ∂K (which can be visualized as the subset of points of ∂K illuminated by a light source at q). Let q' be the closest point of K to q (see Fig 2(b)).

The distance from q to q' is called the *width* of the dual cap. A dual cap of width ε is called an ε -dual cap. The point q is the dual cap's *apex*, and $h(q') \cap \text{conv}(\{q\} \cup K)$ is the dual cap's *base*, which we denote by Γ_q .

Given a surface patch $X \subseteq \partial K$, let $\text{Vor}(X)$ denote the set of points exterior to K whose closest point on ∂K lies within X . Let S denote the $(d-1)$ -dimensional sphere of radius 3 centered at the origin, which we call the *Dudley hypersphere*. (The name is inspired by the sphere used in polytope approximation construction [12]. The value 3 can be replaced by any constant that is at least a constant factor larger than $1 + \varepsilon$.) Throughout, we assume that K has been scaled and translated to lie within the ball of unit radius centered at the origin. It follows that the distance from any point of S to its closest point of K is at least 2. Given a cap or dual cap X , $\text{Vor}(X) \cap S$ is the surface patch consisting of points of S whose closest point on

³The term “dual cap” arises from the observation that, given an inclusion-reversing dual transformation, if we let K^* and q^* be the respective duals of K and q , the dual cap of K generated by q is the dual of the cap of K^* generated by q^* .




Figure 3: Voronoi cells: (a) associated with a surface patch X and (b) a dual cap D_q and the associated Voronoi cell on $H(q)$.

∂K lies within X (see Figure 3(a)).

Because we are interested in the local geometry of K 's boundary, it will be useful to limit the diameter of caps and dual caps. Given $\delta > 0$, define $B_\delta(q)$ to be a closed Euclidean ball of radius δ centered at q . Define the δ -restriction of C_q , denoted C_q^δ , to be $C_q \cap B_\delta(q)$. Similarly, define the δ -restriction of D_q , denoted D_q^δ , to be $D_q \cap B_\delta(q)$.

3 Caps, Dual Caps, and Areas

In this section we present the principal geometric underpinnings of our approximation results. In particular, we consider the relationship between the surface areas of both restricted and unrestricted ε -dual caps and the surface areas of their associated Voronoi patches on S . As mentioned in the introduction, an ε -dual cap of small area must have high curvature, and hence its Voronoi patch on the Dudley hypersphere should be large. We will establish this relationship for dual caps by providing a lower bound on the product of the areas of the ε -dual cap D_q and its associated Voronoi patch, $\text{Vor}(D_q) \cap S$. Analogous results can be shown for caps, but dual caps suffice for our purposes. Our analysis will reveal an interesting similarity between these products and the Mahler volume of a convex body [4, 16, 25]. In Section 3.1 we start with the simpler unrestricted case, and in Section 3.2 we extend this to the restricted case.

3.1 Area Product Bounds: Unrestricted Case

In this section we limit attention to unrestricted dual caps. Rather than dealing with spherical patches, it will be easier to start with hyperplane patches and then generalize to the spherical case. Given a point q external to K , let q' be the closest point on ∂K to q , and recall that $h(q')$ is the unique supporting hyperplane at q' . Let $H(q)$ be a hyperplane that is parallel to $h(q')$, tangent to the Dudley hypersphere S , and on the same side of K as $h(q')$ (see Figure 3(b)). We call this the *Dudley hyperplane* for q . We prove two bounds regarding the

product of the area of dual caps and their Voronoi patches, one for the Dudley hyperplane and one for the Dudley hypersphere.

LEMMA 3.1. *Let $0 < \varepsilon \leq 1$. Let K be a smooth, convex body that lies within a unit ball in \mathbb{R}^d , and let q be a point at distance ε from K . Let D_q denote the ε -dual cap generated by q . Then there is a constant c , depending only on dimension, such that:*

- (i) $\text{area}(D_q) \cdot \text{area}(\text{Vor}(D_q) \cap H(q)) \geq c \cdot \varepsilon^{d-1}$, where $H(q)$ is the Dudley hyperplane for q ,
- (ii) if K is ε -smooth, then $\text{area}(D_q) \cdot \text{area}(\text{Vor}(D_q) \cap S) \geq c \cdot \varepsilon^{d-1}$, where S is the Dudley hypersphere for K .

Before considering $\text{Vor}(D_q) \cap H(q)$, it will be helpful to define a body that is related to the polar dual of the base Γ_q . The parallel hyperplanes $h(q')$ and $H(q)$ can each be identified with \mathbb{R}^{d-1} , by taking the origin to be located at the orthogonal projection of q onto each hyperplane. (In the case of $h(q')$ this is q' .) The coordinate axes for both are taken to be any consistent set of $d-1$ orthonormal vectors lying parallel to these hyperplanes. Since Γ_q is the intersection of a $(d-1)$ -dimensional hyperplane with a bounded convex set that contains q' (the convex hull of $K \cup \{q\}$), we may view it as a convex body in \mathbb{R}^{d-1} that contains the origin. Let $\|v\|$ denote the Euclidean length of $v \in \mathbb{R}^{d-1}$. For each $v \in \Gamma_q$, we define h_v to be a halfspace in \mathbb{R}^{d-1} containing the origin whose bounding hyperplane is orthogonal to v and at distance $\varepsilon/\|v\|$ from the origin (on the same side of the origin as v).

Define $\widehat{\Gamma}_q = \bigcap_{v \in \Gamma_q} h_v$, which we treat as a subset of $H(q)$. This body is bounded, convex, and contains the origin. Rather than analyzing $\text{Vor}(D_q) \cap H(q)$ directly, we will focus on $\widehat{\Gamma}_q$. First, we show that $\widehat{\Gamma}_q$ is a subset of $\text{Vor}(D_q) \cap H(q)$.




Figure 4: Proof of Lemma 3.2.

LEMMA 3.2. *Given the preconditions of Lemma 3.1, $\widehat{\Gamma}_q \subseteq \text{Vor}(D_q) \cap H(q)$.*

Proof. Consider any point $w \in \partial K \setminus D_q$. It suffices to show that w cannot be the closest point of ∂K to any point $p \in \widehat{\Gamma}_q$. It will simplify matters to consider the problem in a 2-dimensional setting. Consider the unique plane Φ containing the points q , q' , and w . For the sake of visualization, we identify this plane with \mathbb{R}^2 so that q' is the origin and the axes are positioned so that q lies on the y -axis at distance ε above q' (see Figure 4(a)). Let K' , $h'(q')$ and $H'(q)$ denote the respective intersections of K , $h(q')$, and $H(q)$ with this plane, and let p' denote the orthogonal projection of p onto Φ . Without loss of generality, we may assume that w lies in the positive- x halfplane. Let t be the point of tangency of $\partial K'$ with respect to q in this halfplane. Let t' denote the point at which the line \overline{qt} intersects $h'(q')$.

Clearly, any point of $\partial K'$ can be closest to some point of $\widehat{\Gamma}_q$ only if it lies on the upper hull of K' (that is, the portion of $\partial K'$ visible from a point at $y = +\infty$). Throughout, we restrict attention to the portion of the upper hull of K' that lies in the positive- x halfplane.

Given a point $u \in \Phi$, let u_x and u_y denote its coordinates relative to the above coordinate system. Let v denote the horizontal vector from q' to t' , whose length is clearly t'_x . By viewing v as a vector in \mathbb{R}^{d-1} on the plane $h(q')$, it contributes a halfspace h_v to the body $\widehat{\Gamma}_q$, which is bounded by a hyperplane that is orthogonal to v (and hence orthogonal to Φ) and lies at distance ε/t'_x from the orthogonal projection of q onto $H'(q)$ (see Figure 4(b)).

To complete the proof, we will show that w is not the closest point on $\partial K'$ to p' . To see why this suffices, observe that the squared Euclidean distance from p to any point on Φ can be decomposed into two

components, one parallel to Φ and one orthogonal to Φ . The orthogonal component is the same for all points on Φ , and so may be ignored when comparing distances. The parallel component is determined by the distances to p' . Thus, if w is not the closest point on $\partial K'$ to p' , then w is not the closest point of ∂K to p .

Since K' is convex and smooth, there is a support-based condition for any point $u \in \partial K'$ to be the closest point on $\partial K'$ to a point $p' \notin K'$. In particular, the line segment $\overline{up'}$ must be perpendicular to the (unique) supporting hyperplane to K' at u . To apply this, consider the convex body $\text{conv}(\{q\} \cup K')$, and let T_q denote the surface patch on this boundary consisting of the union of closed line segments joining q to its points of tangency on K' . (We just show the right-side tangent in Figure 4(b).) Based on the support-based criterion given above, is easy to see that, with respect to the body $\text{conv}(\{q\} \cup K)$, $\text{Vor}(T_q) \cap H(q)$ is the same as $\text{Vor}(D_q) \cap H(q)$. ($\text{Vor}(T_q) \cap H(q)$ is shown as a solid line segment on $H(q)$ in Figure 4(b). Although T_q is not smooth, it can be made so by taking its Minkowski sum with a ball of infinitesimally small radius.)

To show that w is not the closest point on $\partial K'$ to p' , we will show that $\|qp'\| < \|wp'\|$. To this end, consider a ray in Φ emanating from q in the direction of an outward normal to the segment \overline{qt} . The absolute value of the slope of the line \overline{qt} is ε/t'_x , and therefore, the slope of this ray is t'_x/ε . Let ℓ denote the distance from q to $H(q)$. Since each point of K is at distance at least 2 from $H(q)$, we have $\ell \geq 2 - \varepsilon \geq 1$, it follows that this ray intersects $H'(q)$ at a distance of $\ell\varepsilon/t'_x$ from the y -axis, which is to the right of the point at which h_v 's bounding hyperplane intersects $H'(q)$. Since $p \in h_v$, p' lies to the left of this ray. It is easy to see, however, that, restricted to the positive- x halfplane, any point lying to

the left of the ray is closer to q than to any point lying below the line \overline{qt} . Since $w \notin D_q$, it lies below \overline{qt} , and therefore, we achieve the desired conclusion that w is not the closest point on $\partial K'$ to p' . \square

In order to bound the product of the areas of Γ_q and $\widehat{\Gamma}_q$, we will demonstrate the relationship between $\widehat{\Gamma}_q$ and the polar dual of Γ_q , and then apply a well known result from the theory of convexity, called the *Mahler volume*. First, we recall some facts about the *polar transform*. Given a vector $v \in \mathbb{R}^d$ other than the origin, define $\text{polar}(v)$ to be the halfspace that contains the origin whose bounding hyperplane is orthogonal to v and at distance $1/\|v\|$ from the origin (on the same side of the origin as v). Generally, given $r > 0$, define $\text{polar}_r(v)$ to be the halfspace at distance $r^2/\|v\|$ (see Figure 5(a)). Given a convex body K that contains the origin in its interior, define $\text{polar}_r(K)$ to be the intersection of the halfspaces $\text{polar}_r(v)$, for all $v \in K$ (see Figure 5(b)). Clearly, $\text{polar}_r(K)$ is a scaled copy of $\text{polar}(K)$ by a factor of r^2 .




Figure 5: The generalized polar transform and polar body.

The bodies $\text{polar}_{\sqrt{\varepsilon}}(K)$ and $\widehat{\Gamma}_q$ are bounded by the same set of halfspaces, and so we have:

LEMMA 3.3. *Given a smooth convex body K and an exterior point q , $\widehat{\Gamma}_q = \text{polar}_{\sqrt{\varepsilon}}(\Gamma_q)$.*

Given the above results, we can now prove the first part of Lemma 3.1. Recall that $0 < \varepsilon \leq 1$. Since Γ_q is a subset of the orthogonal projection of D_q onto the plane $h(q')$, $\text{area}(\Gamma_q) \leq \text{area}(D_q)$. By Lemmas 3.2 and 3.3, we have

$$\text{area}(\text{Vor}(D_q) \cap H(q)) \geq \text{area}(\widehat{\Gamma}_q) = \text{area}(\text{polar}_{\sqrt{\varepsilon}}(\Gamma_q)).$$

As observed earlier, $\text{polar}_r(\Gamma_q)$ is a scaled copy of $\text{polar}(\Gamma_q)$ by a factor of r^2 . Since this is a $(d-1)$ -dimensional set, we have $\text{area}(\text{polar}_{\sqrt{\varepsilon}}(\Gamma_q)) = \varepsilon^{d-1} \cdot \text{area}(\text{polar}(\Gamma_q))$. By known inequalities on the Mahler volume [16], there is a constant c depending only on

d such that $\text{area}(\Gamma_q) \cdot \text{area}(\text{polar}(\Gamma_q)) \geq c$. Combining these observations, we have

$$\begin{aligned} & \text{area}(D_q) \cdot \text{area}(\text{Vor}(D_q) \cap H(q)) \\ & \geq \text{area}(\Gamma_q) \cdot (\varepsilon^{d-1} \cdot \text{area}(\text{polar}(\Gamma_q))) \\ & \geq c \cdot \varepsilon^{d-1}, \end{aligned}$$

which establishes Lemma 3.1(i).

In order to establish the second part of Lemma 3.1, we need to assume that K is sufficiently smooth. If not, Γ_q might have very small area (much smaller than ε^{d-1}), but the entire surface area of the Dudley hypersphere is bounded above by a constant. If K is ε -smooth, then it is easy to show that Γ_q contains a $(d-1)$ -dimensional ball (centered at q') of radius at least $\varepsilon/\sqrt{3}$.

Consider the area of $\widehat{\Gamma}_q$. Each of its bounding halfspaces h_v is defined by a hyperplane whose distance from the origin is $\varepsilon/\|v\|$. By smoothness, along any direction, there exists a vector $v \in \Gamma_q$ of length at least $\varepsilon/\sqrt{3}$, so the associated bounding hyperplane h_v for $\widehat{\Gamma}_q$ is at distance at most $\varepsilon/(\varepsilon/\sqrt{3}) = \sqrt{3}$ from the origin on $H(q)$. The projection of any differential element on $H(q)$ lying within a constant distance of the origin undergoes a constant factor decrease⁴ when projected onto the Dudley hypersphere. Thus, for some constant factor c' , we have $\text{area}(\text{Vor}(D_q) \cap S) \geq c' \cdot \text{area}(\text{Vor}(D_q) \cap H(q))$. Combining this with Lemma 3.1(i) establishes Lemma 3.1(ii).

3.2 Area Product Bounds: Restricted Case In

this section we generalize the results of the previous section to the case of restricted dual caps. (As before, analogous results hold for restricted caps, but these results are not needed here.) Recall that, given a parameter $\delta > 0$, the δ -restricted cap C_q^δ and the δ -restricted dual cap D_q^δ come about by intersecting the standard cap and dual cap, respectively, with a ball of radius δ centered at q . Restriction complicates the simple relationship between $\widehat{\Gamma}_q$ and the polar body. Although the analysis of the section is structurally similar to that of the previous section, there are added technical complications.

As before, we start by considering the intersection of the Voronoi region with the Dudley hyperplane $H(q)$

⁴Consider a differential element dv on $H(q)$. The corresponding vector with respect to the center of the Dudley hypersphere is of length $x = \sqrt{9 + \|v\|^2}$, and it is at angle $\theta = \arctan(\|v\|/3)$ with respect to vertical. Projecting this element onto the Dudley hypersphere results in a scaling by $(3/x) \cos \theta$. Since $\|v\| \leq \sqrt{3}$, we have $x \leq 2\sqrt{3}$ and $\cos \theta \geq \sqrt{3}/2$. Therefore, the projection of this differential element onto the Dudley sphere is at least $(3/4)dv$, which implies that the projection results in a constant-factor decrease in the surface area.




Figure 6: The definitions of Γ_q^δ and $\widehat{\Gamma}_q^\delta$. (The figure is not drawn to scale.)

and then generalize to the Dudley hypersphere S .

LEMMA 3.4. *Let $0 < \varepsilon \leq 1$ and $3\varepsilon \leq \delta \leq 3\sqrt{\varepsilon}$. Let K be a smooth, convex body that lies within a unit ball in \mathbb{R}^d , and let q be a point at distance ε from K . Let D_q^δ denote the δ -restriction of the ε -dual cap generated by q . Then there is a constant c , depending only on dimension, such that:*

- (i) $\text{area}(D_q^\delta) \cdot \text{area}(\text{Vor}(D_q^\delta) \cap H(q)) \geq c \cdot \delta^{2(d-1)}$, where $H(q)$ is the Dudley hyperplane for q ,
- (ii) if K is ε -smooth, then $\text{area}(D_q^\delta) \cdot \text{area}(\text{Vor}(D_q^\delta) \cap S) \geq c \cdot \delta^{2(d-1)}$, where S is the Dudley hypersphere for K .

Define Γ_q^δ to be the δ -restricted base of the dual cap, by which we mean the intersection of the base with a ball of radius δ centered at q . Because q lies at distance ε above $h'(q)$, Γ_q^δ is the intersection of Γ_q and a $(d-1)$ -dimensional ball of radius $\delta' = \sqrt{\delta^2 - \varepsilon^2}$, centered at q' (see Figure 6(a)). As before, we define an intermediate convex body upon which to base our analysis. Recall that $h(q')$ and $H(q)$ can each be identified with \mathbb{R}^{d-1} by taking the origin to be located at the orthogonal projection of q onto each hyperplane. We may therefore view Γ_q^δ as a convex body in \mathbb{R}^{d-1} that contains the origin. For each $v \in \Gamma_q^\delta$, we define a halfspace h_v in \mathbb{R}^{d-1} as follows. If $\|v\| \geq \delta/3$, then h_v is the halfspace containing the origin whose bounding hyperplane is orthogonal to v and at distance $\delta/6$ from the origin (see Figure 6(b)). If $\|v\| < \delta/3$, then h_v is the $(d-1)$ -dimensional halfspace containing the origin whose bounding hyperplane is orthogonal to v and at distance $\varepsilon/(2\|v\|)$ from the origin (see Figure 6(c)).

Define $\widehat{\Gamma}_q^\delta = \bigcap_{v \in \Gamma_q^\delta} h_v$, which we treat as a subset of $H(q)$. Note that, when $\|v\| < \delta/3$, the hyperplane bounding h_v in this definition is closer to the origin than that given in the definition of $\widehat{\Gamma}_q^\delta$ from the previous section. This body is bounded, convex, and contains the origin. Analogous to Lemma 3.2 we have the following.

LEMMA 3.5. *Given the preconditions of Lemma 3.4, $\widehat{\Gamma}_q^\delta \subseteq \text{Vor}(D_q^\delta) \cap H(q)$.*

Proof. Consider any point $w \in \partial K \setminus D_q^\delta$. If $w \notin D_q$, then the proof follows from the same argument given in Lemma 3.2 (with a suitable adjustment of constants). Thus, it suffices to consider only the case where $w \in D_q \setminus B_\delta(q)$.

As in the earlier proof, define the plane Φ containing the points q , q' , and w , and define K' , $h'(q')$, $H'(q)$, t , and t' as done there. As before, it suffices to show that w is not the closest point on K' to the orthogonal projection of p onto Φ , which we denote by p' . Let v denote the horizontal vector from q' to t' . By viewing v as a vector in \mathbb{R}^{d-1} on the plane $h(q')$, it contributes a halfspace h_v to the body $\widehat{\Gamma}_q^\delta$, which is bounded by a hyperplane that is orthogonal to v (and hence orthogonal to Φ).

We consider two cases, depending on t'_x , the length of v . In the first case, assume that $t'_x \geq \delta/3$ (see Figure 7(a)). We assert that $w_x > \delta/3$. As mentioned earlier, we may assume that $w \notin B_\delta(q)$. Let σ_w and σ_t denote the absolute values of the slopes of the lines $\overline{q'w}$ and \overline{qt} , respectively. Since $w \in D_q$ and K' is convex, the line $\overline{q'w}$ must intersect the line \overline{qt} , which implies that $\sigma_w \leq \sigma_t$. Since $\sigma_t = \varepsilon/t'_x$, and $t'_x \geq \delta/3 \geq \varepsilon$, we have $\sigma_w \leq 1$. Since $w \notin B_\delta(q)$, we have $w_x \geq \delta/\sqrt{2} > \delta/3$, as desired.

Having established that $w_x > \delta/3$, consider the vector v . Since it is of length $t'_x \geq \delta/3$, $\widehat{\Gamma}_q^\delta$ is a subset of the halfspace h_v (in $H(q)$) whose bounding hyperplane is orthogonal to v and at distance $\delta/6$ from the origin (the orthogonal projection of q onto $H(q)$). Since $p' \in \widehat{\Gamma}_q^\delta$, it follows that it lies within distance $\delta/6$ of the y -axis. In contrast, w lies at more than twice this distance from the y -axis. It follows directly that p' is closer to q' than to w , and thus, w cannot be the closest point to p' on $\partial K'$.

In the second case, assume that $t'_x < \delta/3$. Let t''




Figure 7: Proof of Lemma 3.5.

be the intersection of the line \overline{qt} with $\partial B_\delta(q)$, and let q'' be the point on the y -axis horizontal from t'' (see Figure 7(b)). Let σ'' be the absolute slope of the line $\overline{q't''}$. By convexity, and the fact that $w \notin B_\delta(q)$, it follows that w lies below the line $\overline{q't''}$. We assert that $\sigma'' \geq \sigma_t/2$. It is easy to see that this is equivalent to the inequality $q'_y - q''_y \geq (q_y - q''_y)/2$, or equivalently $q_y - q'_y \leq (q_y - q''_y)/2$. By similar triangles, this latter inequality is equivalent to $\|t' - q\| \leq \|t'' - q\|/2$. Clearly, $\|t'' - q\| = \delta$. Since $3\varepsilon \leq \delta$ and $t'_x < \delta/3$, we have

$$\begin{aligned} \|t' - q\| &= \sqrt{t'^2_x + \varepsilon^2} \leq \sqrt{\left(\frac{\delta}{3}\right)^2 + \left(\frac{\delta}{3}\right)^2} \\ &< \frac{\delta}{2} = \frac{\|t'' - q\|}{2}. \end{aligned}$$

This establishes that $\sigma'' \geq \sigma_t/2$, as desired.

Observe that $\sigma_t = \varepsilon/t'_x$. Given that $\sigma'' \geq \sigma_t/2 = \varepsilon/(2t'_x)$, consider a ray emanating from q' that is orthogonal to $\overline{q't''}$. Letting ℓ denote the distance from q to $H(q)$, by similar triangles we see that this ray intersects $H'(q)$ at a distance from the y -axis of $(\ell + \varepsilon)\sigma'' \geq (\ell + \varepsilon)\varepsilon/(2t'_x)$. As observed in the proof of Lemma 3.2, $\ell \geq 1$, and so this ray intersects $H'(q)$ to the right of the point where the hyperplane defining h_v intersects $H'(q)$, whose distance from the y -axis is $\varepsilon/(2\|v\|) = \varepsilon/(2t'_x)$. Since $p \in \widehat{\Gamma}_q^\delta$, it lies to the left of this point. It is easy to see, however, that, restricted to the positive- x halfplane, any point lying to the left of the ray is closer to q' than to any point that lies below the line $\overline{q't''}$. Since w is such a point, we achieve the desired conclusion that w is not the closest point on $\partial K'$ to p' . \square

In order to bound the product of the areas of Γ_q^δ and $\widehat{\Gamma}_q^\delta$, we proceed as in the preceding section by

demonstrating the relationship between $\widehat{\Gamma}_q^\delta$ and the polar dual of Γ_q^δ .

LEMMA 3.6. *Given the preconditions of Lemma 3.4, $\text{polar}_r(\Gamma_q^\delta) \subseteq \widehat{\Gamma}_q^\delta$, where $r \leq \min(\sqrt{\varepsilon/2}, \delta/5)$.*

Proof. Recall that $\text{polar}_r(\Gamma_q^\delta) = \bigcap_{v \in \Gamma_q^\delta} \text{polar}_r(v)$ and $\widehat{\Gamma}_q^\delta = \bigcap_{v \in \Gamma_q^\delta} h_v$, where h_v was defined earlier. For each $v \in \Gamma_q^\delta$, $\text{polar}_r(v)$ and h_v are halfspaces with parallel supporting hyperplanes (both on the same side of the origin), which may differ only with respect to their distances from the origin.

We consider the two cases arising in the definition of h_v . If $\|v\| \geq \delta/3$, then h_v is at distance $\delta/6$ from the origin and $\text{polar}_r(v)$ is at distance $r^2/\|v\| \leq (\delta/5)^2/(\delta/3) \leq \delta/6$, which is at least as close. On the other hand, if $\|v\| < \delta/3$, then h_v is at distance $\varepsilon/(2\|v\|)$ and $\text{polar}_r(v)$ is at distance $r^2/\|v\| \leq \varepsilon/(2\|v\|)$, which is also at least as close. Since each bounding halfspace of $\text{polar}_r(\Gamma_q^\delta)$ is as close to the origin as its parallel bounding halfspace of $\widehat{\Gamma}_q^\delta$, it follows that $\text{polar}_r(\Gamma_q^\delta) \subseteq \widehat{\Gamma}_q^\delta$. \square

Given the above results, we can now prove the first part of Lemma 3.4 in a manner analogous to the proof of Lemma 3.1.

Proof. (of Lemma 3.4)

Recall that Let $0 < \varepsilon \leq 1$ and $\delta \leq 3\sqrt{\varepsilon}$. Since Γ_q^δ is a subset of the orthogonal projection of D_q^δ onto the plane $h(q')$, $\text{area}(\Gamma_q^\delta) \leq \text{area}(D_q^\delta)$. Let $r = \delta/5$. Since $\delta \leq 3\sqrt{\varepsilon}$, $r \leq 3\sqrt{\varepsilon}/5 \leq \sqrt{\varepsilon/2}$. Thus, r satisfies the conditions of Lemma 3.6. By applying both Lemmas 3.5 and 3.6, we have

$$\text{area}(\text{Vor}(D_q^\delta) \cap H(q)) \geq \text{area}(\widehat{\Gamma}_q^\delta) \geq \text{area}(\text{polar}_r(\Gamma_q^\delta)).$$

As observed earlier, $\text{polar}_r(\Gamma_q^\delta)$ is a scaled copy of $\text{polar}(\Gamma_q^\delta)$ by a factor of $r^2 = (\delta/5)^2$. Since this is a $(d-1)$ -dimensional set, we have $\text{area}(\text{polar}_r(\Gamma_q^\delta)) = (\delta/5)^{2(d-1)} \cdot \text{area}(\text{polar}(\Gamma_q^\delta))$. As before, applying the inequality from [16] on the Mahler volume, there is a constant c' depending only on d such that $\text{area}(\Gamma_q^\delta) \cdot \text{area}(\text{polar}(\Gamma_q^\delta)) \geq c'$. Combining these observations, we have

$$\begin{aligned} & \text{area}(D_q^\delta) \cdot \text{area}(\text{Vor}(D_q^\delta) \cap H(q)) \\ & \geq \text{area}(\Gamma_q^\delta) \left(\left(\frac{\delta}{5} \right)^{2(d-1)} \text{area}(\text{polar}(\Gamma_q^\delta)) \right) \\ & \geq \frac{c'}{5^{2(d-1)}} \cdot \delta^{2(d-1)}. \end{aligned}$$

Setting $c = c'/(5^{2(d-1)})$ completes the proof of Lemma 3.4(i). The second part is established by the same argument used in the proof of Lemma 3.1(ii). \square

4 Polytope Approximation

In this section we consider how the results of the previous section⁵ can be applied to polytope approximation. Our main result is presented in the following theorem, which establishes a uniform bound on the complexity of convex approximation based on area, rather than diameter. As mentioned in the introduction, this bound is never significantly worse than Dudley's, but may be significantly better for skinny convex bodies. The proof appears later in this section.

THEOREM 4.1. *Let $K \subset \mathbb{R}^d$ be a convex body such that the width of K in any direction is at least ε . There exists an ε -approximating polytope P with number of facets*

$$O(r \log r), \quad \text{for } r = \frac{\sqrt{\text{area}(K)}}{\varepsilon^{(d-1)/2}}.$$

(This can be adapted to show that there exists an approximating polytope with $O(r \log r)$ vertices instead of facets.) Note that given that an approximating polytope P with $O(r \log r)$ facets exists, then a polytope

⁵The results of Section 3 rely on the assumption that K is ε -smooth. Given a general (non-smooth) convex body, we can apply a simple reduction to the smooth case, as follows. First, we dilate K by performing a Minkowski sum with a Euclidean ball of radius $\varepsilon/2$. Let K' be the result. Observe that the Hausdorff distance between K' and K is $\varepsilon/2$ and K' is $(\varepsilon/2)$ -smooth. Under our assumption that the width of K in any direction is at least ε , this dilation does not increase the surface area of K by more than a constant factor. Next, we apply the results of this section and subsequent approximations, but with the parameter ε set to $\varepsilon/2$. The combination of the two error bounds ($\varepsilon/2$ from dilation and $\varepsilon/2$ from subsequent approximations) results in a total error of ε . The modification of the ε term only affects the constant factors.

with at most $O(r \log^2 r)$ facets can be constructed using [9]. In order to prove the above theorem, we start by stating a lemma that connects dual caps and polytope approximation. The proof follows from the works of Dudley [12] and Bronshteyn and Ivanov [5]. An analogous lemma relates caps and polytope approximation.

LEMMA 4.2. *Let N be any set of points on ∂K that stabs the set of all ε -dual caps of K . Then the polytope defined by tangent planes constructed at the points of N is an ε -approximation of K .*

In order to use ε -nets to stab all dual caps, we need to show that the range spaces of dual caps have constant VC-dimension. Throughout, S denotes the sphere of radius 3 centered at the origin.

LEMMA 4.3. *Let $K \subset \mathbb{R}^d$ be a convex body contained within the unit ball centered at the origin. All the following range spaces (X, R) have VC-dimension at most $d+1$.*

- (1) $X = \partial K$ and R is the set of ε -dual caps.
- (2) $X = S$ and R is the set of Voronoi cells of the ε -dual caps.

Proof. First, we consider an additional range space defined by $X = \partial K$ and R is the set of ε -caps. An ε -cap only contains the points that the corresponding half-space contains, and therefore has VC-dimension upper bounded by the VC-dimension of d -dimensional half-spaces, which is $d+1$. The range space (1) is then equivalent to d -dimensional halfspaces by point-hyperplane duality. To prove the lemma for (2), note that the Voronoi mapping defines a bijection between ∂K and S , and therefore this range space is combinatorially identical to the previous one. \square

Proof. (of Theorem 4.1)

Without loss of generality, we may assume that K has been uniformly scaled to lie within a unit ball centered at the origin and that ε has been scaled accordingly (because the ratio $\text{area}(K)/\varepsilon^{(d-1)}$ is unaffected by uniform scaling). By Lemma 4.2, if we can stab all ε -dual caps of K with $O(r \log r)$ points, then we have an ε -approximating polytope with the same number of facets. We distinguish between two types of ε -dual caps. A dual cap D is of type 1, if $\text{area}(D) \geq \sqrt{\text{area}(K)} \cdot \varepsilon^{(d-1)/2}$, and otherwise it is of type 2.

We use different strategies for stabbing the two types of dual caps. For dual caps of type 1, we use standard machinery to build a $(1/r)$ -net N_1 of size $O(r \log r)$ in the range space (1) of Lemma 4.3. The net N_1 stabs all dual caps of type 1 because of their sufficiently large area.

From Lemma 3.1, for a type-2 dual cap D we have $\text{area}(\text{Vor}(D) \cap S) = \Omega(\varepsilon^{(d-1)/2} / \sqrt{\text{area}(K)})$. In order to stab type-2 dual caps, we build an $O(1/r)$ -net N_2 for the range space (2) of Lemma 4.3. The net N_2 stabs the Voronoi cells of all type-2 dual caps. We then select the nearest neighbor on ∂K for each point of N_2 , obtaining a set N'_2 . The set N'_2 stabs all dual caps of type 2, which concludes the proof. \square

5 Polytope Membership Queries

Consider a convex polytope $K \subset [-1, 1]^d$. In *polytope membership queries*, we are given a query point q , and are to determine whether q lies inside K subject to an error bound ε . Points within distance ε of the boundary of K may be reported as either inside or outside of K . A simple and natural algorithm, called **SplitReduce**, was presented in [2] for preprocessing K in order to answer such queries efficiently. Although an initial analysis was given in [2], the complexity of the resulting data structure remains a challenging problem. In this section, we show that, without increasing storage, it is possible to improve the query time bounds to essentially the square root of those given in [2].

Before describing the algorithm, we need a definition. We say that P is an ε -approximation of K within a quadtree box Q if, for any point $x \in Q$, the following holds: (a) if $x \in K$, then $x \in P$, and (b) if x is at distance at least ε from K , then $x \notin P$. Note that this definition implies that, for any query point $x \in Q$, we can correctly answer ε -approximate polytope membership queries with respect to K by checking whether $x \in P$.

The **SplitReduce** algorithm is presented below. It is given the polytope $K \subset [-1, 1]^d$, the approximation parameter ε , and the desired query time t . The initial quadtree box is $Q = [-1, 1]^d$.

SplitReduce(Q):

1. Let P be an ε -approximation of K within Q .
2. If the number of facets $|P| \leq t$, then Q stores the hyperplanes bounding P .
3. Otherwise, split Q into 2^d quadtree boxes and invoke **SplitReduce** on each such box.

The output of **SplitReduce**(Q) is a quadtree whose leaf cells induce a subdivision of Q . Each leaf cell \widehat{Q} stores a set of halfspaces of size at most t whose intersection approximates K within \widehat{Q} . Let $t(\widehat{Q})$ denote the size of this set. The storage of this quadtree is defined as the total number of halfspaces stored over all the leaf cells. A query is answered by performing a point location in the quadtree followed by a brute force inspection of the halfspaces in the resulting leaf node. Because

cells of diameter ε are not subdivided, the height of the quadtree, and hence the time for point location, is $O(\log(1/\varepsilon))$. For the sake of analysis, let us make the simplifying assumption that the number $t(\widehat{Q})$ of facets computed in Step 1 is minimum. (If instead, Clarkson's approximation algorithm is used [9], the query time increases by a negligible factor of $O(\log(1/\varepsilon))$.)

The main result of this section is the following space-time tradeoff for **SplitReduce**.

THEOREM 5.1. *Given a polytope K of diameter 1, an approximation bound $\varepsilon > 0$ and a real constant $\alpha \geq 2$, it is possible to answer ε -approximate polytope membership queries in time $t = 1/\varepsilon^{(d-1)/(2\alpha)} \log(1/\varepsilon) \geq 1$ from a data structure of space*

$$O\left(\left(\frac{1}{\varepsilon}\right)^{(d-1)\left(1 - \frac{1}{2^{\lfloor \log_2 \alpha \rfloor}} - \frac{\lfloor \log_2 \alpha \rfloor - 1}{\alpha}\right)}\right).$$

In order to establish the tradeoff, we first show how to generalize Lemma 4.2 to the technically more challenging context of approximating K within a box Q .

LEMMA 5.2. *Let K be a convex body (of arbitrary diameter). Let Q be a quadtree box that intersects ∂K . Let E be the portion of ∂K that is within distance $\sqrt{\varepsilon}$ of Q . Let N be any set of points that stab the set of all $\sqrt{\varepsilon}$ -restricted ε -dual caps contained within E , and let P be the polytope defined by the tangent planes constructed at the points of N . Then P is an ε -approximation of K within Q .*

Proof. Let x be any point in Q . Since P is defined by the planes tangent to K , it follows that if $x \in K$ then $x \in P$. Next suppose that x is at distance at least ε from K . We will show that $x \notin P$, which will prove the lemma. Consider the line segment joining x to the point in $K \cap Q$ that is closest to it. By continuity, there must be a point q on this segment that is at distance exactly ε from K . Since Q is convex, this segment and, hence, q must be contained in Q . Consider the $\sqrt{\varepsilon}$ -restricted dual cap whose apex is q . Note that this dual cap is contained within E , and so there is a point of N that stabs it. Clearly, the tangent plane at this point separates q from K , and hence x from K . Thus $x \notin P$, which completes the proof. \square

As in [2], we start by analyzing the low space case. Theorem 5.1 can then be proved by induction. Our key idea is to use the machinery we developed in Section 3. The proof of the following lemma is similar to the proof of Theorem 4.1, but using Lemma 3.4 (with δ set to $\sqrt{\varepsilon}$) instead of Lemma 3.1. Recall that S is the Dudley hypersphere, which we defined as the $(d-1)$ -dimensional sphere of radius 3 centered at the origin.

LEMMA 5.3. Let K, Q and E be as in Lemma 5.2, such that K is ε -smooth and Q is of diameter at least $\sqrt{\varepsilon}$. Further, let E be contained within a ball of unit radius centered at the origin. Then it is possible to ε -approximate K within Q by a polytope with number of facets

$$O(r \log r), \text{ for } r = \frac{\sqrt{\text{area}(E) \cdot \text{area}(\text{Vor}(E) \cap S)}}{\varepsilon^{(d-1)/2}}.$$

Proof. By Lemma 5.2, if we can stab the set of all $\sqrt{\varepsilon}$ -restricted ε -dual caps contained within E with $O(r \log r)$ points, then we have an ε -approximating polytope within Q with $O(r \log r)$ facets. To prove this lemma, we distinguish between two types of such dual caps. A dual cap D is of type 1, if $\text{area}(D) \geq \sqrt{\text{area}(E)/\text{area}(\text{Vor}(E) \cap S)} \cdot \varepsilon^{(d-1)/2}$, and otherwise it is of type 2.

By a straightforward generalization of Lemma 4.3, it follows that both types of restricted dual caps are of constant VC-dimension. We use different strategies for stabbing the two types of dual caps. For dual caps of type 1, we use standard machinery to build a $(1/r)$ -net N_1 , of size $O(r \log r)$ for the range space (1) of Lemma 4.3, considering only the portion E of ∂K . The net N_1 stabs all dual caps of type 1 because of their sufficiently large area.

It follows from Lemma 3.4 that for a type-2 dual cap D , we have

$$\text{area}(\text{Vor}(D) \cap S) = \Omega \left(\varepsilon^{(d-1)/2} \sqrt{\frac{\text{area}(\text{Vor}(E) \cap S)}{\text{area}(E)}} \right).$$

In order to stab type-2 caps, we build an $O(1/r)$ -net N_2 of size $O(r \log r)$ for the range space (2) of Lemma 4.3, considering only $\text{Vor}(E) \cap S$. The net N_2 stabs the Voronoi cell of all type-2 caps. We then select the nearest neighbor on ∂K for each point of N_2 , obtaining a set N'_2 . The set N'_2 stabs all caps of type 2, which concludes the proof. \square

The following lemma will be useful in establishing the storage bounds. Throughout the remainder of the analysis, Q will denote a quadtree box, and T will denote the quadtree produced by $\text{SplitReduce}(Q)$ for some choice of parameters $t \geq 1$ and $\varepsilon < \text{diam}(Q)$. Given T , let T' denote the subtree of T induced by cells of diameter at least $\sqrt{\varepsilon} \text{diam}(Q)/2$. Let L_1 denote the leaf cells of T' that are not subdivided further (i.e., they are leaf cells of T), and let L_2 denote the remaining leaf cells of T' . Observe that if \hat{Q} is any cell of L_2 , then the diameter of \hat{Q} is at most $\sqrt{\varepsilon} \text{diam}(Q)$, since otherwise its children would have diameter more than $\sqrt{\varepsilon} \text{diam}(Q)/2$, and \hat{Q} would not be a leaf cell of T' .

LEMMA 5.4.

- (i) $\sum_{\hat{Q} \in L_1} t(\hat{Q}) = O((\text{diam}(Q)/\varepsilon)^{(d-1)/2})$.
- (ii) $|L_2| = O \left(\left(\frac{\log t}{t} \right)^2 \cdot \left(\frac{\text{diam}(Q)}{\varepsilon} \right)^{(d-1)/2} \right)$.

Proof. The proof of (i) follows from the proof of Lemma 2.2 in [2]. For completeness, we provide a short sketch. Dudley's approximation yields an ε -approximation to $K \cap Q$ as the intersection of $O((\text{diam}(Q)/\varepsilon)^{(d-1)/2})$ halfspaces. In Dudley's construction, each of these halfspaces is responsible for approximating a region of K 's boundary of diameter $O(\sqrt{\varepsilon} \text{diam}(Q))$. By a simple packing argument, each halfspace is useful in this manner for only a constant number of cells $\hat{Q} \in L_1$, which establishes assertion (i).

To prove (ii), let \hat{Q} be any cell of L_2 . Recall that $\sqrt{\varepsilon} \text{diam}(Q)/2 \leq \text{diam}(\hat{Q}) \leq \sqrt{\varepsilon} \text{diam}(Q)$. Let $\hat{\varepsilon} = \varepsilon/(4 \text{diam}(Q))$. For the remainder of this proof, we take the center of Q as the origin and assume that we have scaled space so that Q has diameter 1. Note that in this scaled space $\sqrt{\hat{\varepsilon}} \leq \text{diam}(\hat{Q}) \leq 2\sqrt{\hat{\varepsilon}}$. As usual, we let S denote the Dudley hypersphere, defined as the $(d-1)$ -dimensional sphere of radius 3 centered at the origin. Let \hat{K} be the Minkowski sum of K and a ball of radius $\hat{\varepsilon}$ (thus, \hat{K} is $\hat{\varepsilon}$ -smooth). Let \hat{E} denote the portion of $\partial \hat{K}$ that is within distance at most $\sqrt{\hat{\varepsilon}}$ of \hat{Q} . Note that \hat{E} is contained within the unit ball centered at the origin.

The intuition underlying our approach for bounding the number of cells in L_2 is to show that $\text{area}(\text{Vor}(\hat{E}) \cap S)$ must be large. Since the area of S is a constant, we will show that this implies that there cannot be too many cells in L_2 . In order to show that $\text{area}(\text{Vor}(\hat{E}) \cap S)$ is large, we will apply Lemma 5.3 and show that otherwise we could have approximated K within \hat{Q} to the desired accuracy with fewer than t facets and \hat{Q} would then not have been subdivided further. We now make these ideas precise.

We first prove that $\text{area}(\text{Vor}(\hat{E}) \cap S) = \Omega((t/\log t)^2 \cdot \hat{\varepsilon}^{(d-1)/2})$. For the sake of contradiction, suppose that $\text{area}(\text{Vor}(\hat{E}) \cap S) \leq c \cdot (t/\log t)^2 \cdot \hat{\varepsilon}^{(d-1)/2}$ for a suitable constant c . Clearly, the conditions of Lemma 5.3 are satisfied with $\hat{K}, \hat{Q}, \hat{E}$, and $\hat{\varepsilon}$ in place of K, Q, E , and ε , respectively. Therefore, \hat{K} can be $\hat{\varepsilon}$ -approximated within \hat{Q} by a polytope with m facets, where

$$m \leq c_1 r \log r, \text{ for } r = \frac{\sqrt{\text{area}(\hat{E}) \cdot \text{area}(\text{Vor}(\hat{E}) \cap S)}}{\hat{\varepsilon}^{(d-1)/2}},$$

where c_1 is the constant hidden in the O -notation of Lemma 5.3. Note that an $\hat{\varepsilon}$ absolute error in the space after scaling corresponds to $\hat{\varepsilon} \cdot \text{diam}(Q) = \varepsilon/4$ absolute

error prior to scaling. Thus, prior to scaling, m facets suffice to provide an ε -approximation of K within \widehat{Q} . To compute m , observe that

$$\text{area}(\widehat{E}) \leq c_2 \text{diam}(\widehat{Q})^{d-1} \leq c_2 2^{d-1} \widehat{\varepsilon}^{(d-1)/2},$$

where c_2 is some constant. By our assumption, $\text{area}(\text{Vor}(\widehat{E}) \cap S) \leq c \cdot (t/\log t)^2 \cdot \widehat{\varepsilon}^{(d-1)/2}$. Substituting these upper bounds on areas, we obtain

$$\begin{aligned} r &\leq \frac{\sqrt{c_2 2^{d-1} \widehat{\varepsilon}^{(d-1)/2} \cdot c (t/\log t)^2 \widehat{\varepsilon}^{(d-1)/2}}}{\widehat{\varepsilon}^{(d-1)/2}} \\ &\leq \sqrt{2^{d-1} c_2 c} \cdot \frac{t}{\log t}. \end{aligned}$$

Recalling that $m \leq c_1 r \log r$, it is easy to see that we can make $m < t$ by choosing c to be sufficiently small. Therefore, \widehat{Q} should not have been subdivided further, which contradicts the fact that $\widehat{Q} \in L_2$. Thus $\text{area}(\text{Vor}(\widehat{E}) \cap S) = \Omega((t/\log t)^2 \cdot \widehat{\varepsilon}^{(d-1)/2}) = \Omega((t/\log t)^2 \cdot (\varepsilon/\text{diam}(\widehat{Q}))^{(d-1)/2})$.

Since the area of S is a constant and any point of S can be contained in $\text{Vor}(\widehat{E})$ for at most a constant number of cells $\widehat{Q} \in L_2$, it follows that the cardinality of L_2 is $O((\log t/t)^2 \cdot (\text{diam}(\widehat{Q})/\varepsilon)^{(d-1)/2})$, which completes the proof. \square

We are now ready to present our improved analysis of the low space case.

LEMMA 5.5. *The output of $\text{SplitReduce}(Q)$ for*

$$t \geq \left(\frac{\text{diam}(Q)}{\varepsilon} \right)^{(d-1)/8} \log \left(\frac{\text{diam}(Q)}{\varepsilon} \right) \geq 1$$

is a quadtree with storage $O((\text{diam}(Q)/\varepsilon)^{(d-1)/2})$.

Proof. Let T denote the quadtree produced by the algorithm. We will show that $\sum t(\widehat{Q})$ over all leaf cells \widehat{Q} of T is $O((\text{diam}(Q)/\varepsilon)^{(d-1)/2})$, which will establish the desired storage bound. Recall T', L_1 , and L_2 as defined just before Lemma 5.4. By Lemma 5.4(i), $\sum_{\widehat{Q} \in L_1} t(\widehat{Q}) = O((\text{diam}(Q)/\varepsilon)^{(d-1)/2})$.

Next, consider the set L_2 of cells. Let \widehat{Q} be any cell of L_2 . Recall that $\sqrt{\varepsilon} \text{diam}(\widehat{Q})/2 \leq \text{diam}(\widehat{Q}) \leq \sqrt{\varepsilon} \text{diam}(Q)$. Since $t \geq (\text{diam}(Q)/\varepsilon)^{(d-1)/8}$, it follows that $t \geq (\text{diam}(\widehat{Q})/\varepsilon)^{(d-1)/4} \geq 1$ for $\varepsilon \leq \text{diam}(Q)/4$ (if $\varepsilon > \text{diam}(Q)/4$ then the lemma holds trivially).

Lemma 2.2 in [2] states that the output of $\text{SplitReduce}(\widehat{Q})$ for $t \geq (\text{diam}(\widehat{Q})/\varepsilon)^{(d-1)/4} \geq 1$ is a quadtree with storage $O((\text{diam}(\widehat{Q})/\varepsilon)^{(d-1)/2})$. Therefore, the storage for the subtree rooted at \widehat{Q} is at most $O((\text{diam}(\widehat{Q})/\varepsilon)^{(d-1)/2}) = O((\text{diam}(Q)/\varepsilon)^{(d-1)/4})$. Also, by Lemma 5.4(ii), $|L_2| = O((\log t/t)^2 \cdot$

$(\text{diam}(Q)/\varepsilon)^{(d-1)/2}) = O((\text{diam}(Q)/\varepsilon)^{(d-1)/4})$, where we have used the lower bound on t given in the statement of the lemma. It follows that the total storage for the subtrees rooted at cells of L_2 is at most $O((\text{diam}(Q)/\varepsilon)^{(d-1)/4} \cdot (\text{diam}(Q)/\varepsilon)^{(d-1)/4}) = O((\text{diam}(Q)/\varepsilon)^{(d-1)/2})$. This completes the proof. \square

Using Lemma 5.5 as the base case, we now extend the space-time tradeoff to other query times.

THEOREM 5.6. *Let $\alpha \geq 2$ be a constant. The output of $\text{SplitReduce}(Q)$ for*

$$t \geq \left(\frac{\text{diam}(Q)}{\varepsilon} \right)^{(d-1)/(2\alpha)} \log \left(\frac{\text{diam}(Q)}{\varepsilon} \right) \geq 1$$

is a quadtree with storage

$$O \left(\left(\frac{\text{diam}(Q)}{\varepsilon} \right)^{(d-1) \left(1 - \frac{1}{2^{\lfloor \log_2 \alpha \rfloor}} - \frac{\lfloor \log_2 \alpha \rfloor - 1}{\alpha} \right)} \right).$$

Proof. Let $k = \lfloor \log_2 \alpha \rfloor$. Our proof proceeds by induction on a constant number of steps k . The base case $k = 1$ corresponds to Lemma 5.5. Next, assume that the theorem holds for $1, \dots, k-1$, that is, for $\alpha < 2^k$. We need to prove that the theorem holds for $2^k \leq \alpha < 2^{k+1}$.

Let T denote the quadtree produced by the algorithm with $2^k \leq \alpha < 2^{k+1}$. Recall T', L_1 , and L_2 as defined just before Lemma 5.4. By Lemma 5.4(i), $\sum_{\widehat{Q} \in L_1} t(\widehat{Q}) = O((\text{diam}(Q)/\varepsilon)^{(d-1)/2})$, and thus the storage for the leaf cells of L_1 satisfies the bound of the theorem.

We next consider the set L_2 of cells. Let \widehat{Q} be any cell of L_2 . Recall that $\sqrt{\varepsilon} \text{diam}(\widehat{Q})/2 \leq \text{diam}(\widehat{Q}) \leq \sqrt{\varepsilon} \text{diam}(Q)$. It follows that

$$t \geq \left(\frac{\text{diam}(\widehat{Q})}{\varepsilon} \right)^{(d-1)/(2\alpha')} \log \left(\frac{\text{diam}(\widehat{Q})}{\varepsilon} \right) \geq 1$$

where $\alpha' = \alpha/2$ for $\varepsilon \leq \text{diam}(Q)/16$ (if $\varepsilon > \text{diam}(Q)/16$, then the theorem holds trivially). Therefore, we can use the induction hypothesis to obtain the following storage for each cell $\widehat{Q} \in L_2$:

$$\begin{aligned} &O \left(\left(\frac{\text{diam}(\widehat{Q})}{\varepsilon} \right)^{(d-1) \left(1 - \frac{1}{2^{k-1}} - \frac{k-2}{\alpha'} \right)} \right) = \\ &O \left(\left(\frac{\text{diam}(Q)}{\varepsilon} \right)^{(d-1) \left(\frac{1}{2} - \frac{1}{2^k} - \frac{k-2}{\alpha} \right)} \right). \end{aligned}$$

Also, by Lemma 5.4(ii), $|L_2| = O((\log t/t)^2 \cdot (\text{diam}(Q)/\varepsilon)^{(d-1)/2}) = O((\text{diam}(Q)/\varepsilon)^{(d-1) \left(\frac{1}{2} - \frac{1}{\alpha} \right)})$,

where we have used the value of t given in the statement of the theorem. It follows that the total storage for the subtrees rooted at cells of L_2 is at most

$$O\left(\left(\frac{\text{diam}(Q)}{\varepsilon}\right)^{(d-1)\left(1-\frac{1}{2^k}-\frac{k-1}{\alpha}\right)}\right).$$

This completes the proof. \square

Recalling that the time for locating the leaf cell containing the query point is $O(\log(1/\varepsilon))$, Theorem 5.1 follows as an immediate consequence.

References

- [1] P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. Geometric approximation via coresets. In J. E. Goodman, J. Pach, and E. Welzl, editors, *Combinatorial and Computational Geometry*. MSRI Publications, 2005.
- [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, T. Malamatos, and D. M. Mount. Space-time tradeoffs for approximate nearest neighbor searching. *J. Assoc. Comput. Mach.*, 57:1–54, 2009.
- [4] J. Bourgain and V. D. Milman. New volume ratio properties for convex symmetric bodies. *Inventiones Mathematicae*, 88:319–340, 1987.
- [5] E. M. Bronshteyn and L. D. Ivanov. The approximation of convex sets by polyhedra. *Siberian Math. J.*, 16:852–853, 1976.
- [6] E. M. Bronstein. Approximation of convex sets by polytopes. *Journal of Mathematical Sciences*, 153(6):727–762, 2008.
- [7] T. M. Chan. Fixed-dimensional linear programming queries made easy. In *Proc. 12th Annu. Sympos. Comput. Geom.*, pages 284–290, 1996.
- [8] T. M. Chan. Optimal partition trees. In *Proc. 26th Annu. Sympos. Comput. Geom.*, pages 1–10, 2010.
- [9] K. L. Clarkson. Algorithms for polytope covering and approximation. In *Proc. 3th Workshop Algorithms Data Struct.*, pages 246–252, 1993.
- [10] K. L. Clarkson. Building triangulations using epsilon-nets. In *Proc. 38th Annu. ACM Sympos. Theory Comput.*, pages 326–335, 2006.
- [11] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 3rd edition, 2010.
- [12] R. M. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *Approx. Theory*, 10(3):227–236, 1974.
- [13] R. Fleischer, K. Mehlhorn, G. Rote, E. Welzl, and C.-K. Yap. On simultaneous inner and outer approximation of shapes. In *Proc. Sixth Annu. Sympos. Comput. Geom.*, pages 216–224, 1990.
- [14] P. M. Gruber. Asymptotic estimates for best and stepwise approximation of convex bodies I. *Forum Math.*, 5:521–537, 1993.
- [15] J. K. Boroczky. Approximation of general smooth convex bodies. *Adv. Math.*, 153:325–341, 2000.
- [16] G. Kuperberg. From the Mahler conjecture to Gauss linking integrals. *Geometric And Functional Analysis*, 18:870–892, 2008.
- [17] M. A. Lopez and S. Reisner. Efficient approximation of convex polygons. *Internat. J. Comput. Geom. Appl.*, 10:445–452, 2000.
- [18] J. Matoušek and O. Schwarzkopf. On ray shooting in convex polytopes. *Discrete Comput. Geom.*, 10:215–232, 1993.
- [19] J. Matoušek. Reporting points in halfspaces. *Comput. Geom. Theory Appl.*, 2:169–186, 1992.
- [20] J. Matoušek. Linear optimization queries. *J. Algorithms*, 14(3):432–448, 1993.
- [21] D. E. McClure and R. A. Vitalie. Polygonal approximation of plane convex bodies. *J. Math. Anal. Appl.*, 51:326–358, 1975.
- [22] J. S. B. Mitchell and S. Suri. Separation and approximation of polyhedral objects. *Comput. Geom. Theory Appl.*, 5:95–114, 1995.
- [23] E. A. Ramos. Linear programming queries revisited. In *Proc. 16th Annu. Sympos. Comput. Geom.*, pages 176–181, 2000.
- [24] G. Rote. The convergence rate of the sandwich algorithm for approximating convex functions. *Computing*, 48:337–361, 1992.
- [25] L. A. Santaló. An affine invariant for convex bodies of n -dimensional space. *Portugaliae Mathematica*, 8:155–161, 1949. (In Spanish).
- [26] R. Schneider. Polyhedral approximation of smooth convex bodies. *J. Math. Anal. Appl.*, 128:470–474, 1987.
- [27] L. F. Toth. Approximation by polygons and polyhedra. *Bull. Amer. Math. Soc.*, 54:431–438, 1948.