Optimal Area-Sensitive Bounds for Polytope Approximation

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

To appear at ACM Symposium on Computational Geometry (SoCG), 2012.

Abstract

Approximating convex bodies is a fundamental question in geometry and has applications to a wide variety of optimization problems. Given a convex body K in $\mathbb{R}^d$ for fixed 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. The best known uniform bound, due to Dudley (1974), shows that $O((\mathrm{diam}(K)/\varepsilon)^{(d-1)/2})$ facets suffice. While this bound is optimal in the case of a Euclidean ball, it is far from optimal for skinny convex bodies.

We show that, under the assumption that the width of the body in any direction is at least ε, it is possible to approximate a convex body using $O(\sqrt{\mathrm{area}(K)}/\varepsilon^{(d-1)/2})$ facets, where $\mathrm{area}(K)$ is the surface area of the body. This bound is never worse than the previous bound and may be significantly better for skinny bodies. This bound is provably optimal in the worst case and improves upon our earlier result (which appeared in SODA 2012).

Our improved bound arises from a novel approach to sampling points on the boundary of a convex body in order to stab all (dual) caps of a given width. This approach involves the application of an elegant concept from the theory of convex bodies, called *Macbeath regions*. While Macbeath regions are defined in terms of volume considerations, we show that by applying them to both the original body and its dual, and then combining this with known bounds on the Mahler volume, it is possible to achieve the desired width-based sampling.

*This author's work was supported by the Research Grants Council, Hong Kong, China under project number 610108.

†Research supported by CNPq and FAPERJ grants.

‡This author's work was supported in part by the National Science Foundation under grant CCF-1117259 and the Office of Naval Research under grant N00014-08-1-1015.

1 Introduction

Approximating convex bodies by polytopes is a fundamental problem, which has been extensively studied in the literature. (See Bronstein [13] for a recent survey.) At issue is the minimum number of vertices (alternatively, the minimum number of facets) needed in an approximating polytope for a given error $\varepsilon > 0$. Consider a convex body K in Euclidean d-dimensional space. A polytope P is said to *ε-approximate* K if the Hausdorff distance [13] between K and P is at most ε. Throughout, we will restrict attention to the Hausdorff metric, and we assume that the dimension d is a constant.

Our interest is in establishing bounds on the combinatorial complexity of approximating general convex bodies. Approximation bounds are of two common types. In both cases, it is shown that there exists $\varepsilon_0 > 0$ such that 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 the combinatorial complexity of the approximating polytope approaches infinity. Examples include bounds by Gruber [22], Clarkson [16], and others [9, 26, 28, 29].

In the second type, which we call *uniform bounds*, the value of ε_0 is independent of K. For example, these include the results of Dudley [19] and Bronshteyn and Ivanov [12]. These bounds hold without any smoothness assumptions. Dudley showed that, for $\varepsilon \leq 1$, any convex body K can be ε-approximated by a polytope P with $O((\mathrm{diam}(K)/\varepsilon)^{(d-1)/2})$ facets. Bronshteyn and Ivanov showed the same bound holds for the number of vertices. Constants hidden in the O-notation depend only on d. These results have many applications, for example, in the construction of coresets [1].

The approximation bounds of both Dudley and Bronshteyn and Ivanov are tight up to constant factors (specifically when K is a Euclidean ball). These bounds may be significantly suboptimal if K is skinny, however. In an earlier paper [3], we presented an upper bound that is based not on diameter, but on surface area. In particular, let $\mathrm{area}(K)$ denote the $(d-1)$-dimensional Hausdorff measure of ∂K. We showed that, under the assumption that the width of the body in any direction is at least ε, there exists an ε-approximating polytope whose number of facets is $O(t \log t)$, where $t = \sqrt{\mathrm{area}(K)}/\varepsilon^{(d-1)/2}$. For a given diameter, the surface area of a convex body is maximized for a Euclidean ball, implying that $\mathrm{area}(K) = O(\mathrm{diam}(K)^{d-1})$. Thus, this bound is tight in the worst case up to the logarithmic factor. The additional log factor is disconcerting since it implies that the bound is suboptimal even for the simple case of a Euclidean ball. In this paper we show that the logarithmic factor can be eliminated. In particular, we prove the following result, which is worst-case optimal, up to constant factors.

Theorem 1.1. *Consider real d-space, $\mathbb{R}^d$. There exists a positive ε_0 and constant c_d such that for any convex body $K \subset \mathbb{R}^d$ and any ε, $0 < \varepsilon \leq \varepsilon_0$, if the width of K in any direction is at least ε, then there exists an ε-approximating polytope P whose number of facets is at most*

$$c_d \sqrt{\mathrm{area}(K)}/\varepsilon^{(d-1)/2}.$$

Note that the width assumption seems to be a technical necessity. For example, consider a $(d-2)$-dimensional unit ball B embedded within $\mathbb{R}^d$, and let B' denote its Minkowski sum with a d-dimensional ball of radius $\delta \ll \varepsilon$. By the optimality of Dudley's bound for Euclidean balls, $\Omega(1/\varepsilon^{(d-3)/2})$ facets are needed to approximate B and hence to approximate B'. But, the surface area of B' can be made arbitrarily small as a function of δ.

The width assumption is not a fundamental impediment, however. If the body is of width less than ε in some direction, then by projecting the body onto a hyperplane orthogonal to this direction, it is possible to reduce the problem to a convex approximation problem in one lower dimension. This can be repeated until the body's width is sufficiently large in all remaining dimensions, and the stated bound can be applied in this lower dimensional subspace.

1.1 Overview of Methods

It is well known (see, e.g., [11, 12, 15]) that computing a Hausdorff approximation of a convex body K by a polytope can be reduced to the problem of computing an economical set cover or an economical hitting set in which the set systems involve appropriately defined surface patches on K's boundary. Depending on the nature of the approximation (e.g., whether minimizing the number of vertices or facets and whether an inner or outer approximation is desired) the surface patches of interest are either *caps* or *dual caps*. For our purposes, a cap is the portion of K's boundary that lies within a halfspace and a dual cap is the portion of K's boundary that is visible from an external point. (Formal definitions are given in Section 2.4.) For both caps and dual caps, we define a hyperplanar surface, which we call a *base*, whose area is less than or equal to the associate cap or dual cap, respectively. Of particular interest are caps and dual caps whose defining width is ε. Depending on the formulation, the approximation problem reduces to computing a small set of points on the boundary of K such that every cap of width ε contains one of these points or every dual cap of width ε contains one of these points.

The source of slackness in the bound of [3] arises from a sampling method that is based on a relatively heavy-handed tool, namely ε-nets for halfspace ranges. In light of recent lower bounds on the size of ε-nets for halfspace ranges [27], it is clear that the elimination of the log factor requires a sampling process that is specially tailored to caps or dual caps. The principal contribution of this paper is such a sampling method.

Our new approach makes use of a classical structure from the theory of convexity, called *Macbeath regions*. Intuitively, for any convex body K and a volume parameter v, there exists a collection of pairwise disjoint bodies, each of volume $\Omega(v)$, such that for every halfspace H where the cap $K \cap H$ has volume v, one of these bodies will be completely contained within this cap. (The formal statement is given in Section 2.4.) Macbeath regions have found numerous uses in the theory of convex sets and the geometry of numbers (see Bárány [7]). To date, the application of Macbeath regions in the field of computational geometry has been quite limited. For example, they have been used as a technical device in proving lower bounds for range searching (see, e.g., [4, 5, 10]).

Because their definition is based on volume, not width, the use of Macbeath regions in the context of uniform bounds for convex approximation has been limited to volume-based notions of distance, such as the Nikodym metric (which is based on the volume of the symmetric difference) [6, 8]. The difficulty in adapting Macbeath regions to width-based sampling is that caps of a given volume may have widely varying widths. Our approach to dealing with this is through the application of a two-pronged strategy, which combines Macbeath-based sampling in both the original body and its dual.

This strategy relies on a combination of two well known dual transformations, the polar dual (which is based on distances to the origin) and the functional dual (which is the dual transform most widely used in computational geometry and is based on vertical distances). The problem with either form of dual is that distances are not generally preserved, and this makes it difficult to relate approximations in the primal and dual settings. The principal feature of the functional dual transform is that it does preserve vertical distances between points and hyperplanes.

To exploit this, we decompose the approximation problem into a constant number of subproblems, where each involves approximating just a portion of the body in which the surface normals have similar directions. After a suitable rotation, within each subproblem, the distance from an external point to the boundary can be approximated by its vertical distance. In

Lemma 4.1 in Section 4, we establish an intriguing relationship between the base of a dual cap in the original body K and the base of the corresponding cap in the functional dual set K^*. In particular, we show that the base of a dual cap in K is essentially the polar of the base of the corresponding cap in K^*. This polar relationship is an essential ingredient in our construction, because it allows us to apply the classical concept of the *Mahler volume*, to show that small dual caps in K correspond to large caps in K^*. Our two-pronged sampling works because dual caps that are too small to be sampled in K will be sampled as caps in K^*.

2 Preliminaries

Let K denote a convex body in $\mathbb{R}^d$, that is, a compact convex subset with nonempty interior. Let ∂K denote its boundary. Throughout, we assume that the dimension d is a constant. We say that a convex body is *smooth*, if at each point $q \in \partial K$, there exists a ball of positive radius that lies entirely within K and contains q on its boundary.

If K is smooth, there is a unique supporting hyperplane at every point on its boundary. Since we do not assume smoothness, we augment our representation of boundary points. Given any convex body K, we define an *augmented point* on ∂K to consist of a pair (q, u), where q is a point on ∂K and u is an outward-directed unit vector that is orthogonal to some supporting hyperplane passing through q. We call such a vector a *surface normal.* To keep our notation simple, we will usually refer to an augmented point simply as q, but it is understood throughout that every augmented point is associated with a unique surface normal and hence a unique supporting hyperplane, which we will denote by $h(q)$. Define $H^+(q)$ to be the closed halfspace bounded by $h(q)$ that contains K, and define $H^-(q)$ to be the other closed halfspace bounded by $h(q)$.

Given a convex body K in $\mathbb{R}^d$, let $\mathrm{vol}(K)$ denote its d-dimensional Hausdorff measure. Given a $(d-1)$-dimensional manifold Ψ in $\mathbb{R}^d$, for example, a surface patch on a convex body, let $\mathrm{area}(\Psi)$ denote its $(d-1)$-dimensional Hausdorff measure. We use $\mathrm{area}(K)$ as a convenient shorthand for $\mathrm{area}(\partial K)$.

For $\varepsilon > 0$, we say that a polytope P is an *ε-approximation* of K if the Hausdorff distance between K and P is at most ε. Observe that simultaneously scaling K and ε by any positive factor, does not alter the ratio $\sqrt{\mathrm{area}(K)}/\varepsilon^{(d-1)/2}$. Therefore, for the sake of proving Theorem 1.1, we may assume that K has been uniformly scaled to lie within the hypercube $[-(1-2\varepsilon), (1-2\varepsilon)]^d$. This means that the Minkowski sum of K with a ball of radius 2ε (a shape that will be useful to us later) lies entirely within $[-1, 1]^d$.

To avoid specifying many real-valued constants that arise in our analysis, we will often hide them using asymptotic notation. For positive real x, we use the notation $O(x)$ (resp., $\Omega(x)$) to mean a quantity whose value is at most (resp., at least) cx for an appropriately chosen constant c. We use $\Theta(x)$ to denote a quantity that lies within the interval $[cx, c'x]$ for appropriate constants c and c'. These constants will generally be functions of d, but not of ε.

2.1 Nonuniform Area-Based Bounds

Before presenting our analysis, we note that a nonuniform bound very similar to ours can be derived from Gruber's result [22]. (We thank Quentin Merigót for pointing this out.) Gruber shows that if K is a strictly convex body and ∂K is twice differentiable, then as ε approaches zero, the number of bounding halfspaces needed to achieve an ε-approximation of K is

$$O\left(\left(\frac{1}{\varepsilon}\right)^{(d-1)/2} \int_{\partial K} \kappa(x)^{1/2} dx\right), \tag{1}$$

where $\kappa(x)$ denotes the Gaussian curvature of K at x, and dx is a differential surface element. (Böröczky showed that the requirement that K be "strictly" convex can be eliminated [9].) Because square root is concave and $\int_{\partial K} dx = \text{area}(K)$, we may apply Jensen's inequality to obtain

$$\frac{1}{\text{area}(K)} \int_{\partial K} \kappa(x)^{1/2} dx \;\leq\; \left(\frac{1}{\text{area}(K)} \int_{\partial K} \kappa(x) dx \right)^{1/2}.$$

By the Gauss-Bonnet theorem [18], the total Gaussian curvature of K is bounded, from which we conclude that the number of approximating halfspaces is

$$O\left(\frac{\sqrt{\text{area}(K)}}{\varepsilon^{(d-1)/2}} \right),$$

which matches the bound of Theorem 1.1.

We hasten to add that this approach cannot be used to produce a uniform bound, however. To show this, we will present a two-dimensional counterexample (but the result can be readily extended to any constant dimension). Consider a fixed value of ε, and let $0 < \delta \leq \varepsilon$. Let $m = \lfloor 1/\sqrt{\delta} \rfloor$, and define K_δ to be the Minkowski sum of a regular m-gon inscribed in a unit circle and a Euclidean ball of radius δ. Observe that K_δ consists of m flat edges, each of length $\Theta(\sqrt{\varepsilon})$, connected by m circular arcs, each of radius δ and subtending an angle of $2\pi/m$. Since $\delta \leq \varepsilon$, it is straightforward to show that any convex polygon that ε-approximates K_δ requires $\Omega(1/\sqrt{\varepsilon})$ sides. (This follows from the same argument that shows that Dudley's bound is tight for circles.) We also assert that

$$\int_{\partial K_\delta} \kappa(x)^{1/2} dx \;=\; \Theta\left(\sqrt{\delta}\right).$$

To see this, observe that the flat sides of K_δ contribute zero to the integral. Each circular arc has curvature $1/\delta$, and so altogether they contribute a total of $2\pi\delta/\sqrt{\delta} = 2\pi\sqrt{\delta}$ to the integral.

Thus, Equation (1) provides a bound[1] on the number of sides of an ε-approximating polygon of $O\big(\sqrt{\delta/\varepsilon}\big)$, which is clearly incorrect given our hypothesis that $\delta \leq \varepsilon$. Indeed, this hypothesis is exactly the sort of assumption that the nonuniform analysis forbids. In contrast, since Theorem 1.1 is a uniform bound, it can be applied to this type of example.

2.2 Support Sets

Our analysis makes use of a representation of the body K as the intersection of $2d$ unbounded convex sets, based on the orientations of surface normals. Let I denote the index set $\{1, \ldots, d\}$, and let $I^{\pm}$ denote the $(2d)$-element index set consisting of $\{\pm 1, \ldots, \pm d\}$. For $i \in I^{\pm}$, if $i > 0$, let x_i denote the unit vector associated with the ith coordinate axis, and if $i < 0$, let x_i be its negation. Let X_i denote the hyperplane passing through the origin and orthogonal to x_i. (For the sake of illustration, we will think of x_i as being directed vertically upwards and X_i as horizontal.) The *vertical projection* of a point onto X_i means setting its ith coordinate to zero.

For any vector $\vec{u} = (u_1, \ldots, u_d) \in \mathbb{R}^d$, let $\|\vec{u}\|_\infty = \max_i |u_i|$. For $i \in I^{\pm}$ we say that i is a *signed principal axis* of $\vec{u}$ if $\text{sign}(i) \cdot u_i = \|\vec{u}\|_\infty$. (For example, the vector $(-2, 1, 2, -2, 0)$ has the signed principal axes -1, 3, and -4.) For each $i \in I^{\pm}$, let Ψ_i denote the subset of augmented

[1] Note that we cannot apply Gruber's or Böröczky's theorems directly to K_δ, since its boundary is not twice differentiable. In particular, the second derivative is discontinuous at the joints where each edge meets a circular arc. We can easily fix this by creating a sufficiently small gap at each joint and introducing a smooth polynomial spline of constant degree to fill the gap. Although the resulting body is not strictly convex, Böröczky showed that this assumption is not necessary for the bound to hold.

points $q \in \partial K$ such that i is a signed principal axis of the inward-directed surface normal at q. These subsets subdivide K's boundary into $2d$ (relatively closed) surface patches. For each $i \in I^{\pm}$, let $S_i = \bigcap_{q \in \Psi_i} H^+(q)$ (see Figure 1(a)). We call this the ith *support set* of K. Each S_i is an unbounded convex set that properly contains K. Clearly, $\bigcap_{i \in I^{\pm}} S_i = K$.

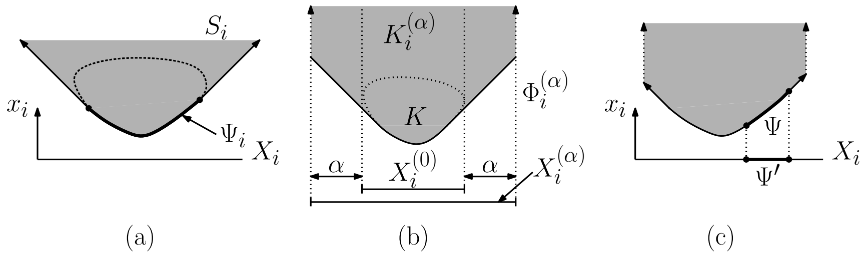

Figure 1: Support sets and projections.

It will be useful to further restrict each support set S_i so that its vertical projection covers only a bounded region of X_i that is somewhat larger than the vertical projection of K. For a given $\alpha \geq 0$, take the Minkowski sum of K with a ball of radius α and project this set vertically onto X_i. Let $X_i^{(\alpha)}$ denote the resulting $(d-1)$-dimensional convex body on X_i. Let $\Phi_i^{(\alpha)}$ denote the infinite cylinder whose central axis is aligned with x_i and whose horizontal cross section is $X_i^{(\alpha)}$ (see Figure 1(b)). Define $K_i^{(\alpha)}$ to be $S_i \cap \Phi_i^{(\alpha)}$, which we call the *$\alpha$-restricted support set.* As we shall see below, the body $K_i^{(2\varepsilon)}$, will play an important role in our analysis. Indeed, to simplify notation, henceforth we let K_i denote $K_i^{(2\varepsilon)}$. By our initial scaling of K, it follows that the vertical projection of K_i lies within the hypercube $[-1, 1]^{d-1}$.

A closed convex set is *U-shaped* (with respect to a given vertical direction) if for every point on the set's boundary, the vertical ray directed upwards from this point lies entirely within the set. When dealing with U-shaped sets, such as K_i, we will be most interested in the lower hull, not the vertical sides. For this reason, given a U-shaped set U, we define ∂U to consist of the augmented points of the boundary of U such that the associated surface normal for each point is not horizontal. In fact, we will impose the stronger restriction that the absolute tangent of the angle between the surface normal and the vertical axis is at most a constant (whose exact value is implicit in the radius of the ball B of Lemma 4.2 below). Define $\text{area}(U)$ to be the surface area of ∂U.

Given a parameter $\sigma > 0$, a U-shaped set U is *σ-steep* if for every augmented point $q \in \partial U$, the absolute tangent of the angle between the vertical axis and the surface normal at q does not exceed σ. The halfspace above a horizontal hyperplane is 0-steep, and K_i is easily seen to be $O(1)$-steep. The following are straightforward consequences of the above definitions.

Lemma 2.1.

(*i*) *If the width of K in any direction is at least ε, then for any $i \in I^{\pm}$,* $\text{area}(K_i) = O(\text{area}(K))$.

(*ii*) *Consider any surface patch Ψ of an $O(1)$-steep U-shaped set, and let Ψ' denote its vertical projection. Then* $\text{area}(\Psi') = \Theta(\text{area}(\Psi))$ *(see Figure 1(c)).*

2.3 Dual Transforms

Our results are based on two commonly used dual transforms in geometry. Such transforms map points to hyperplanes and vice versa, while preserving point-hyperplane incidences. The first transform is a generalization of the standard polar transform, and the second is a dual transform frequently used in computational geometry, which we call the functional dual.

Given a vector $v \in \mathbb{R}^d$ other than the origin and $r > 0$, define the *generalized polar transform* of v, denoted $\mathrm{polar}_r(v)$, to be the halfspace containing the origin, whose bounding hyperplane is orthogonal to v and is at distance $r^2/\|v\|$ from the origin (see Figure 2(a)). More formally, $\mathrm{polar}_r(v) = \{u \in \mathbb{R}^d : u_1v_1 + \cdots + u_dv_d \leq r^2\}$. Given a convex body K that contains the origin in its interior, define $\mathrm{polar}_r(K) = \bigcap_{v\in K} \mathrm{polar}_r(v)$ (see Figure 2(b)).

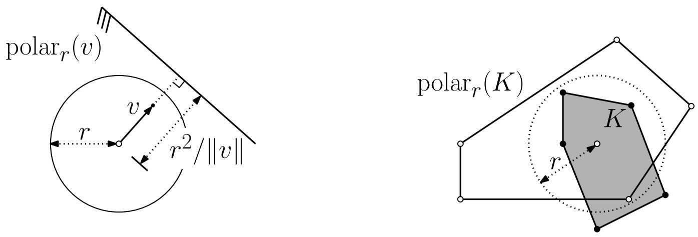

(a) (b)

Figure 2: The generalized polar transform.

The standard geometric polar transform [20] arises as a special case when $r = 1$. In particular, $\mathrm{polar}_r(K)$ is a scaled copy of $\mathrm{polar}_1(K)$ by a factor of r^2.

Later, we will make use of an important result from the theory of convex sets, which states that, given a convex body K, the product $\mathrm{vol}(K) \cdot \mathrm{vol}(\mathrm{polar}_1(K))$, which is called K's *Mahler volume*, is bounded below by a constant (see, e.g., [24]).

Next, let us define the *functional dual transform* [14]. Let $x_1, \ldots, x_{d-1}$ denote the first $d-1$ coordinates of a point, and let y denote the d-th coordinate. Any nonvertical hyperplane $y = \sum_{j=1}^{d-1} a_jx_j - a_d$ can be represented by a d-tuple $(a_1, \ldots, a_d)$. Given a point $p = (p_1, \ldots, p_d) \in \mathbb{R}^d$, define its functional dual, denoted p^*, to be the non-vertical hyperplane:

$$p^* : y = \sum_{j=1}^{d-1} p_jx_j - p_d.$$

(The term "functional" comes from the fact that y is expressed as a function of the x_j's.) The dual of a nonvertical hyperplane is defined so that $p^{**} = p$. When applying this transform in the context of support sets, such as K_i, the ith coordinate axis will take on the role of the dth (vertical) axis in the above definitions.

It is easily verified that the functional dual transform negates vertical distances, in the sense that the signed vertical distance from a point p to a hyperplane h is equal to the negation of the signed vertical distance from p^* to h^* [17]. (The fact that point-hyperplane incidences are preserved is a direct consequence.) Given any U-shaped convex set U, define its dual U^* to be the intersection of the upper halfspaces of q^* for all $q \in \partial U$. It is easily verified that the dual is also U-shaped. Further, if U's vertical projection lies within a ball of constant radius centered at the origin, then U^* is $O(1)$-steep.

Let U be a U-shaped convex body and let U^* be its dual. There exists a natural correspondence between the augmented points on ∂U and augmented points on ∂U^*. Given an augmented point (q, u) on ∂U, let h be the supporting plane at q that is orthogonal to u. Let $p = h^*$. Clearly, p lies on ∂U^*, and the dual hyperplane q^* is a supporting hyperplane at p. Letting v denote the outward-directed unit vector orthogonal to q^*, we define the augmented point on ∂U^* *corresponding* to (q, u) to be (p, v). By the involutory nature of the dual, the augmented point on ∂U corresponding to (p, v) is (q, u).

2.4 Caps, Dual Caps, and Macbeath Regions

Bronshteyn and Ivanov [12] and Dudley [19] demonstrated the relevance of caps and dual-caps (defined below) to convex approximation. Let U be a U-shaped convex set, and let $\varepsilon > 0$. For any augmented point $q \in \partial U$, recall that $h(q)$ is the supporting hyperplane at this point, and let $h(q) + \varepsilon$ denote its translate vertically upwards by ε. The intersection of ∂U and the lower halfspace of $h(q) + \varepsilon$ is called the *ε-cap* induced by q, which we denote by $C_q(U)$ (see Figure 3(a)). (In contrast to standard usage, where a cap consists of the subset of the body lying in the halfspace, for us a cap is a subset of the nonvertical boundary.) The intersection of $h(q) + \varepsilon$ with U is called the *base* of the cap, and is denoted by $\Gamma_q(U)$. When U is clear from context, we simply write C_q and Γ_q.

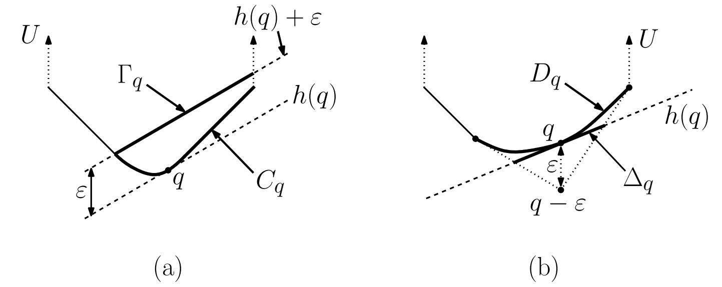

Figure 3: Caps and dual caps.

For any augmented point $q \in \partial U$, let $q - \varepsilon$ be the vertical translate of q downwards by distance ε. Define the *ε-dual cap*, denoted $D_q(U)$, to be the portion of ∂U that is visible from $q - \varepsilon$ (see Figure 3(b)). The intersection of the bounding halfspaces of U that contain $q - \varepsilon$ defines an infinite cone. The intersection of this cone with $h(q)$ is called the dual cap's *base*, and is denoted $\Delta_q(U)$. Again, when U is clear, we simply write D_q and Δ_q.

Given an augmented point $q \in \partial U$, consider cap $C_q(U)$ and dual cap $D_q(U)$. We say that an augmented point $p \in \partial U$ *stabs* $C_q(U)$ if $p \in C_q(U)$, and we say that p *stabs* $D_q(U)$ if the hyperplane $h(p)$ passing through p separates U from the point $q - \varepsilon$. The importance of ε-dual caps to approximation is established in the following lemma. We say that an ε-dual cap of K_i is *useful*, if its inducing point q lies on $\partial K_i^{(\varepsilon)}$. Recall that $K_i = K_i^{(2\varepsilon)}$, so the inducing point is at horizontal distance at least ε from the vertical portion of K_i.

Lemma 2.2. *Let K be a convex body in $\mathbb{R}^d$. For any $\varepsilon > 0$ and $i \in I^{\pm}$, let Q_i be a set of augmented points on ∂K_i that stab all useful ε-dual caps of K_i. Let $P_i = \bigcap_{q \in Q_i} H^+(q)$, and $P = \bigcap_{i \in I^{\pm}} P_i$. Then P is an ε-approximation to K. The number of facets in P is at most $\sum_{i \in I^{\pm}} |Q_i|$.*

Proof. Observe that P is the intersection of supporting halfspaces for K, and hence $K \subseteq P$. Thus, every point of K is within distance 0 of P. To show that every point of P is within distance ε of K, we will prove the contrapositive. In particular, we will show that any point $p \in \mathbb{R}^d$ whose distance from ∂K exceeds ε cannot be in P.

Let q be the closest point of ∂K to p. Let p' be a point along the segment $\overline{pq}$ whose distance from q is exactly ε (see Figure 4). Let i be any signed principal axis for the inward surface normal for K at q. Clearly, the horizontal component of the distance between p' and q is at most ε.

The vertical projection of q lies within $X_i^{(0)}$. Therefore the vertical projection of p' lies within $X_i^{(\varepsilon)}$, which is subset of $X_i^{(2\varepsilon)}$, above which K_i lies. This implies that that there exists a point $q' \in \partial K_i^{(\varepsilon)} \subset \partial K_i$ that is vertically above p' (see Figure 4). Augment q' by associating it

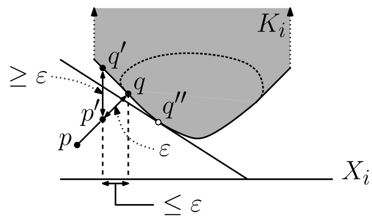

Figure 4: Proof of Lemma 2.2.

with any valid surface normal for which i is a signed principal axis. It follows that any ε-dual cap induced by q' is useful. By local minimality considerations, it is easy to see that q is the closest point to p' on ∂K_i. Therefore, the distance from p' to q' is at least ε. By hypothesis, the ε-dual cap induced by q' is stabbed by the supporting hyperplane of some augmented point $q'' \in Q_i$. This implies that p' lies outside this bounding hyperplane, and hence is external to P_i. It follows directly that p is also external to P_i, and hence it is external to P. This completes the proof. ☐

In order to prove the results of the next two sections, we will make use of the following result. It demonstrates the existence of a collection of convex bodies, such that all caps of sufficiently large volume contain at least one such body. These bodies are closely related to the concept of *Macbeath regions* (also called *M-regions*). This concept was introduced by Macbeath [25], and its relevance to cap coverings was explored by Ewald, Larman, and Rogers [21]. Applications of Macbeath regions have appeared in numerous works, include Bárány and Larman [8], Bárány [7], and Brönnimann *et al.* [10]. The following lemma follows directly from these earlier works.

Lemma 2.3. *Given a convex body $K \subset \mathbb{R}^d$ and a parameter $0 < v \leq \mathrm{vol}(K)$, there exist two collections of convex bodies, $\mathcal{M}$ and $\mathcal{M}'$, such that the bodies of $\mathcal{M}$ are contained within K and are pairwise disjoint. Each $M \in \mathcal{M}$ is associated with a corresponding body in $\mathcal{M}'$, denoted M', such that $M \subseteq M'$. M' is called M's* expanded body. *These sets satisfy the following:*

- (*i*) *For all $M \in \mathcal{M}$, $\mathrm{vol}(M)$ and $\mathrm{vol}(M')$ are both $\Theta(v)$.*
- (*ii*) *For any halfspace H, if $\mathrm{vol}(K \cap H) = v$, then there exists $M \in \mathcal{M}$ such that $M \subseteq K \cap H \subseteq M'$, where M' is M's expanded body.*

3 Stabbing Dual Caps in the Primal

The purpose of this section is to establish bounds on the size of stabbing sets for ε-dual caps. Our approach will involve the use of the Macbeath region machinery from Lemma 2.3. We begin with the following utility lemma, which establishes a relationship between the areas of the bases of the cap and dual cap induced by the same point. Recall that, given an augmented point q on the boundary of a U-shaped set, D_q denotes the ε-dual cap induced by q, and Δ_q is its base. Also, C_q denotes the ε-cap induced by q, and Γ_q denotes its base.

Lemma 3.1. *Let U be any $O(1)$-steep U-shaped convex set in $\mathbb{R}^d$, and let $\varepsilon > 0$. Consider any augmented point $q \in \partial U$ such that the vertical projection of C_q is bounded. Let h be the supporting hyperplane at q, and let $h + \varepsilon$ denote the hyperplane containing Γ_q. Let H be the lower halfspace bounded by $h + \varepsilon$. Then $U \cap H$ contains a (generalized) cone whose base is a vertical translate of Δ_q onto $h + \varepsilon$, and whose apex is at vertical distance ε from the base. Further,* $\mathrm{vol}(U \cap H) = \Theta(\varepsilon \cdot \mathrm{area}(\Delta_q))$.

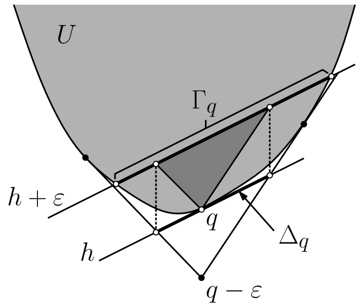

Figure 5: Proof of Lemma 3.1.

Proof. Recall that $q-\varepsilon$ denotes the point at distance ε vertically below q, which forms the apex of D_q (see Figure 5).

Consider the cone T bounded by the supporting hyperplanes of the dual cap and h. The apex of this cone is $q-\varepsilon$, and its base is Δ_q. Let T' denote the vertical translate of T upwards by distance ε. It is easy to see that T' lies entirely within $U \cap H$, which establishes the first claim.

Next, consider the cone T'' bounded by the same supporting hyperplanes as T, but whose base is $h+\varepsilon$. Because T'' is bounded by the supporting hyperplanes for U, it follows that Γ_q is contained within the base of this cone. Clearly, T'' is just a factor-2 scaling of T. Therefore, their volumes are related to one another by a constant factor. Since U is $O(1)$-steep, it follows that $\mathrm{vol}(T') = \mathrm{vol}(T) = \Theta(\varepsilon \cdot \mathrm{area}(\Delta_q))$. It is easy to see that $T' \subseteq U \cap H \subseteq T''$, and therefore $\mathrm{vol}(U \cap H) = \Theta(\varepsilon \cdot \mathrm{area}(\Delta_q))$, which establishes the second claim. □

Our next lemma is the main result of this section. It states that there exists a small set of points that stab all sufficiently large ε-dual caps of K_i. The notion of "large" is based on the area of the dual cap's base and a threshold parameter t.

Lemma 3.2. *Let K be a convex body in $\mathbb{R}^d$, and consider any $\varepsilon > 0$ and $i \in I^{\pm}$. Given any $t > 0$, there exists a set of $O(\mathrm{area}(K)/t)$ augmented points $Q_i \subset \partial K_i$ such that, for any useful ε-dual cap D_q induced by a point $q \in \partial K_i^{(\varepsilon)}$, if $\mathrm{area}(\Delta_q) \geq t$, there exists a point in Q_i that stabs D_q.*

Proof. Apply Lemma 2.3 to K_i with $v = c \cdot \varepsilon \cdot t$, for a suitable constant c, and let $\mathcal{M}$ and $\mathcal{M}'$ denote the resulting collections of bodies. (Note that K_i is not bounded, but for our purposes it suffices to bound it crudely, say by intersection with the lower halfspace of a horizontal hyperplane that is high enough to contain every point at vertical distance ε above ∂K_i.)

For each $M \in \mathcal{M}$, if M does not lie entirely within vertical distance ε of ∂K_i, then discard M from further consideration. (Such a body cannot lie within any useful ε-cap, and hence will be of no value to us.) Otherwise, let $M' \in \mathcal{M}'$ denote M's expanded body as described in Lemma 2.3(ii).

Let N be a net [2] chosen so that, for a suitable constant c', any ellipsoid contained within M' of volume at least $c'v$ contains at least one point of the net. By Lemma 2.3(i), $\mathrm{vol}(M') = \Theta(v)$, so any ellipsoid of volume $\Omega(v)$ covers a constant fraction of the volume of M'. Since the range space of ellipsoids is known to have constant VC-dimension, it follows that the size of the resulting net is $O(1)$. For each point of the net that lies within K_i, project it vertically downward onto ∂K_i, and augment it by associating it with any valid surface normal whose signed principal axis is i. Add the resulting augmented point to Q_i. Repeating this process for all $M \in \mathcal{M}$ yields the desired set Q_i.

To establish the correctness of this construction, consider any augmented point $q \in \partial K_i$ whose ε-dual cap (D_q) has base (Δ_q) of area at least t. Let h denote the supporting hyperplane at q, and let H denote the lower halfspace bounded by $h + \varepsilon$, which contains the base of q's cap, C_q. Clearly, the vertical projection of C_q is bounded, and so by Lemma 3.1, $\mathrm{vol}(K_i \cap H) = \Omega(\varepsilon \cdot \mathrm{area}(\Delta_q)) = \Omega(\varepsilon t)$. By choosing c suitably, we can ensure that $\mathrm{vol}(K_i \cap H) \geq v$. In order to apply Lemma 2.3, observe that $\mathrm{vol}(K_i \cap H)$ is a continuous monotonic function of ε, which decreases to 0 when $\varepsilon = 0$. Therefore, if $\mathrm{vol}(K_i \cap H) > v$, there exists a value $\varepsilon' < \varepsilon$ such that $\mathrm{vol}(K_i \cap H) = v$. The resulting ε'-dual cap and ε'-cap are subsets of the original dual cap and cap, respectively, and therefore there is no loss of generality in using these reduced objects throughout the rest of the analysis in place of the originals.

By Lemma 2.3(ii), there exists $M \in \mathcal{M}$ such that $M \subseteq K_i \cap H \subseteq M'$. (Note that M could not have been discarded in the construction process.) By Lemma 3.1, $K_i \cap H$ contains a cone T whose base is a translate of Δ_q and whose vertical height is ε. This cone is a convex body, and thus by John's Theorem [23], it contains an ellipsoid E such that

$$\mathrm{vol}(E) \;=\; \Omega(\mathrm{vol}(T)) \;=\; \Omega(\varepsilon \cdot \mathrm{area}(\Delta_q)) \;=\; \Omega(v).$$

Therefore, by choosing the constant c' in the net construction appropriately, there exists a point of the net that lies within T, and hence the vertical projection of this point will be included in Q_i. It is easy to see that the horizontal extents of a dual cap's base lie entirely within the horizontal extents of the dual cap itself. Therefore, this point of Q_i stabs the dual cap, as desired.

Finally, we bound the number of points in Q_i. To do this, consider the set of points of K_i that lie within vertical distance ε of its boundary. The volume of this region is clearly $O(\varepsilon \cdot \mathrm{area}(K_i))$, which by Lemma 2.1(i) is $O(\varepsilon \cdot \mathrm{area}(K))$. By Lemma 2.3(i), each body M is of volume $\Omega(v) = \Omega(\varepsilon t)$. These bodies are pairwise disjoint, and (after discarding) each of them lies within vertical distance ε of ∂K_i. Therefore, by a simple packing argument, the number of such bodies is

$$O\left(\frac{\varepsilon \cdot \mathrm{area}(K_i)}{v}\right) \;=\; O\left(\frac{\varepsilon \cdot \mathrm{area}(K)}{\varepsilon \cdot t}\right) \;=\; O\left(\frac{\mathrm{area}(K)}{t}\right),$$

as desired. □

4 Stabbing Caps in the Dual

In this section we will consider the problem of stabbing caps of the (functional) dual of the convex body (recall Section 2.3). Consider a U-shaped set U and let U^* be its dual. The following lemma establishes a polar relationship between the bases of the ε-dual caps of U and the bases of the ε-caps in the dual U^*.

Consider any augmented point $q \in \partial U$ and let p be the corresponding augmented point on ∂U^* (recall the definition from the end of Section 2.3). The lemma shows that the base of q's dual cap and the base of p's cap, if viewed as convex bodies in $\mathbb{R}^{d-1}$, are general polar duals of each other. Given a convex body K and a point p, let $K - p$ denote the translate of K so that p coincides with the origin.

Lemma 4.1. *Let q be an augmented point of ∂U, and let p be the corresponding augmented point on ∂U^*. Let q' and p' be the respective vertical projections of q and p. Let Δ_q be the base of q's ε-dual cap in U, and let Γ_p be the base of p's ε-cap in U^*. Let Δ'_q and Γ'_p denote the respective vertical projections of these bases. Then, $\Delta'_q - q' = \mathrm{polar}_r(\Gamma'_p - p')$, for $r = \sqrt{\varepsilon}$.*

Proof. We begin by showing that Δ_q consists of the set of points $v \in h(q)$ such that, for all $u \in \Gamma_p$,

$$\sum_{j=1}^{d-1} (u_j - p_j)(v_j - q_j) \ \le \ \varepsilon.$$

By the incidence-preserving property of the dual transform, the supporting hyperplane passing through p is q^*. From basic properties of the dual transform, the hyperplane containing Γ_p is the vertical translate of q^* upwards by a distance of ε, which is just $(q-\varepsilon)^*$ (see Figure 6).

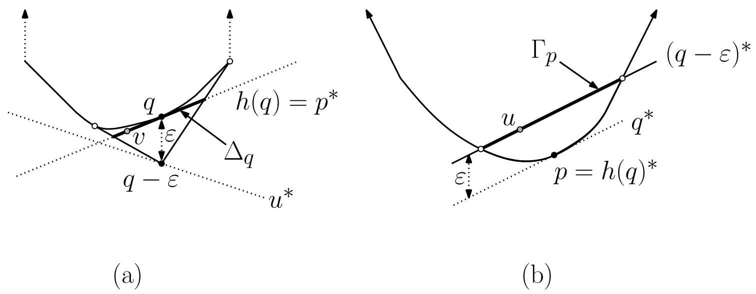

Figure 6: Proof of Lemma 4.1.

The base Δ_q of the ε-dual cap induced by q is the intersection of the supporting hyperplane $h(q)$ with $\mathrm{conv}(\{q-\varepsilon\} \cup U)$. Another way to express this set is to define an infinite cone C formed by the intersection all the upper halfspaces of the hyperplanes that pass through $q-\varepsilon$ such that U lies within each upper halfspace. Clearly, $\Delta_q = C \cap h(q) = C \cap p^*$. By the incidence-preserving property of the dual transform, these hyperplanes are the duals of points u lying on $(q-\varepsilon)^*$. By the order-reversing property of the dual, each such point u lies within U^* (since it lies above all the supporting hyperplanes of U^*). But, the set of points $u \in (q-\varepsilon)^* \cap U^*$ is easily seen to be Γ_p (see the point u in Figure 6(b) and hyperplane u^* in Figure 6(a)). Letting $(u^*)^+$ denote the upper halfspace of u's dual, we have $\Delta_q = \left(\bigcap_{u \in \Gamma_p} (u^*)^+\right) \cap p^*$.

The equations of the hyperplanes q^* and $(q-\varepsilon)^*$ are

$$\begin{aligned} q^* \quad &: \quad y = \sum_{j=1}^{d-1} q_j x_j - q_d \\ (q-\varepsilon)^* \quad &: \quad y = \sum_{j=1}^{d-1} q_j x_j - q_d + \varepsilon. \end{aligned}$$

Since p lies on q^*, we have $p_d = \sum_{j=1}^{d-1} q_j p_j - q_d$. Any $u \in \Gamma_p$ lies on $(q-\varepsilon)^*$, and so we have $u_d = \sum_{j=1}^{d-1} q_j u_j - q_d + \varepsilon$.

From the remarks made earlier, a point v is in Δ_q if and only if: (1) it lies on p^*, and (2) it lies in the upper halfspace $(u^*)^+$ for each $u \in \Gamma_p$. From condition (1) we have

$$\begin{aligned} v_d \ &= \ \sum_{j=1}^{d-1} p_j v_j - p_d \ = \ \sum_{j=1}^{d-1} p_j v_j - \left(\sum_{j=1}^{d-1} q_j p_j - q_d \right) \\ &= \ \sum_{j=1}^{d-1} p_j (v_j - q_j) + q_d. \end{aligned}$$

From condition (2) we have

$$
\begin{aligned}
v_d &\geq \sum_{j=1}^{d-1} u_j v_j - u_d = \sum_{j=1}^{d-1} u_j v_j - \left(\sum_{j=1}^{d-1} q_j u_j - q_d + \varepsilon \right) \\
&= \sum_{j=1}^{d-1} u_j (v_j - q_j) + q_d - \varepsilon.
\end{aligned}
$$

Combining these observations, we have $v \in \Delta_q$ if and only if $v \in p^*$ and for all $u \in \Gamma_p$,

$$
\sum_{j=1}^{d-1} p_j (v_j - q_j) + q_d \geq \sum_{j=1}^{d-1} u_j (v_j - q_j) + q_d - \varepsilon,
$$

or equivalently, $\sum_{j=1}^{d-1} (u_j - p_j)(v_j - q_j) \leq \varepsilon$, as desired.

Recall that Δ'_q and Γ'_p are the respective vertical projections of Δ_q and Γ_p. Recall that p' and q' are the respective vertical projections of p and q. Then, we have shown that $v \in \Delta'_q$ if and only if, for all $u \in \Gamma'_p$, $\sum_{j=1}^{d-1} (u_j - p_j)(v_j - q_j) \leq \varepsilon$. If we translate Δ'_q and Γ'_p so that q' and p' each coincide with the origin, then these bodies are generalized polars of each other, for $r = \sqrt{\varepsilon}$. That is, $\Delta'_q - q' = \mathrm{polar}_r(\Gamma'_p - p')$. □

We will apply the above result to K_i to convert the problem of stabbing ε-dual caps to the problem of stabbing ε-caps in K_i^*. Although we have restricted K_i so that the nonvertical elements of its boundary are bounded, the same cannot be said for its dual K_i^*, which (due to the vertical sides of K_i) has a vertical projection that covers the entire horizontal coordinate hyperplane X_i. Therefore, it would be meaningless to apply the Macbeath region approach that was used in the previous section. What rescues us is the observation that we are only interested in ε-caps of K_i^* that arise from the application of the previous lemma to useful ε-dual caps of K_i.

Recall that each augmented point $q \in \partial K_i^{(\varepsilon)}$ induces a useful ε-dual cap. We say that a cap of K_i^* is *useful* if its inducing augmented point corresponds to an augmented point of $\partial K_i^{(\varepsilon)}$. The next lemma shows that each such cap is of constant horizontal extents.

Lemma 4.2. *Let C_p be any useful ε-cap induced by an augmented point $p \in \partial K_i^*$. Then, the vertical projection of C_p lies entirely within a ball B of radius $O(1)$ centered at the origin.*

Proof. Recall that Γ_p denotes the base of p's cap in K_i^*, and let Γ'_p be its vertical projection. Let q be p's corresponding point. Since p is useful, we have $q \in \partial K_i^{(\varepsilon)}$. Also, recall that Δ_q denotes the base of q's ε-dual cap, and let Δ'_q be its vertical projection. By Lemma 4.1, $\Delta'_q - q' = \mathrm{polar}_r(\Gamma'_p - p')$, for $r = \sqrt{\varepsilon}$, which by the involutory nature of the polar transform implies that $\Gamma'_p - p' = \mathrm{polar}_r(\Delta'_q - q')$.

We assert that there exists a positive constant c such that Δ'_q contains a Euclidean ball centered at q' of radius $c\varepsilon$ (see Figure 7(a)). Assuming this assertion for now, observe that, by basic properties of the generalized polar transform, $\mathrm{polar}_r(\Delta'_q - q')$ is contained within a ball of radius $r^2/c\varepsilon = 1/c$ centered at the origin. Since $\Gamma'_p - p' = \mathrm{polar}_r(\Delta'_q - q')$, it follows that Γ'_p is contained within a ball of radius $1/c$ centered at the p' (see Figure 7(b)). Since K_i is $O(1)$-steep, the coordinates of p' (which are all slopes) are of constant absolute values. Therefore, p' is within constant distance of the origin, and hence so are all the points of Γ'_p.

To complete the proof, it suffices to establish the above assertion. We begin by showing that the cone with apex $q - \varepsilon$ that defines D_q contains a cone $\Psi(q)$ whose apex is at $q - \varepsilon$ and whose

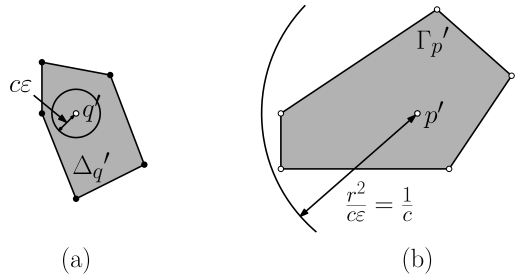

Figure 7: Proof of Lemma 4.2.

central angle is $\Omega(1)$ (see Figure 8(a)). To see this, consider any augmented point $u \in \partial K_i$ that lies on the boundary of Δ_q. This means that, the supporting hyperplane at u passes through $q - \varepsilon$. It suffices to show that the acute angle between the ray from $q - \varepsilon$ to u and the vertical axis is $\Omega(1)$.

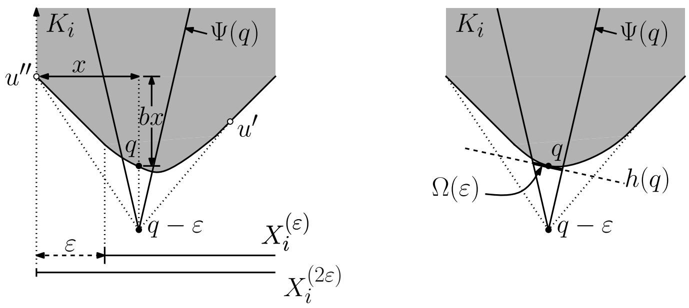

Figure 8: Proof of Lemma 4.2 (continued).

If u is a point in the relative interior of K_i (such as u' in Figure 8(a)), then, by the slope constraints of the associated support set, the angle of the ray from $q - \varepsilon$ to u and the vertical axis is $\Omega(1)$. On the other hand, suppose that u is a point on the relative boundary of K_i (such as u'' in Figure 8(a)). Since C_p is useful, it follows that D_q is useful, which implies that $q \in \partial K_i^{(\varepsilon)}$.

Recall that K_i is the intersection of S_i with the infinite cylinder $\Phi_i^{(2\varepsilon)}$ whose cross section is $X_i^{(2\varepsilon)}$. By definition, $X_i^{(2\varepsilon)}$ contains a region of distance at least ε around the vertical projection of any $q \in \partial K_i^{(\varepsilon)}$. If we let x denote the horizontal component of the distance from q to u, we have $x \geq \varepsilon$. By the constraint on the slopes of K_i's support set, since both q and u lie on ∂K_i, the vertical component of the distance between q and u is at most bx, for some positive constant b. Since $q - \varepsilon$ lies directly below q by a distance of ε, the horizontal component of the distance between $q - \varepsilon$ and u is x, and the vertical component of the distance is at most $bx + \varepsilon \leq (b+1)x$. Therefore, the tangent of angle between the ray from $q - \varepsilon$ to u and the vertical axis is at least $x/(b+1)x = 1/(b+1) = \Omega(1)$. Therefore, for a suitably chosen (constant) central angle, the cone $\Psi(q)$ lies within the cone defining D_q.

Next, consider the supporting hyperplane $h(q)$ that passes through the augmented point q. Because $q \in \partial K_i^{(\varepsilon)}$, this hyperplane satisfies the slope constraints of the associated support set. Because q is located at distance ε above $\Psi(q)$'s apex and $h(q)$ is of constant slope, it follows that $h(q) \cap \Psi(q)$ is an ellipse whose smallest radius is $\Omega(\varepsilon)$ (see Figure 8(b)). Therefore, the vertical projection of this ellipse contains a ball of radius $\Omega(\varepsilon)$. Because $\Psi(q)$ lies within the cone defining D_q, it follows that this ball lies entirely within the vertical projection Δ_q'. This establishes the assertion, and completes the proof. $\square$

The following lemma establishes a cap-based variant of the sampling process of Lemma 3.2. The proof is a straightforward adaptation of the proof of Lemma 3.2 to the case of caps. Caps are actually easier to deal with because the Macbeath-region machinery can be applied directly.

Lemma 4.3. *Let K be a convex body in $\mathbb{R}^d$, and consider any $\varepsilon > 0$ and $i \in I^{\pm}$. Given any $t > 0$, there exists a set of $O(1/t)$ augmented points $P_i \subset \partial K_i^*$ such that, for any useful ε-cap C_p induced by an augmented point $p \in \partial K_i^*$, if $\mathrm{area}(\Gamma_p) \geq t$, there exists a point in P_i that stabs C_p.*

Proof. Apply Lemma 2.3 to K_i^* with $v = c \cdot \varepsilon \cdot t$, for a suitable constant c, and let $\mathcal{M}$ denote the resulting collection of disjoint bodies of volume $\Omega(v)$. (Note that K_i^* is not bounded, but it suffices to bound it crudely, say by a ball of suitably large radius.) Let B denote the $(d-1)$-dimensional ball of radius $O(1)$ that lies on X_i, as described in Lemma 4.2.

As in the proof of Lemma 3.2, for each $M \in \mathcal{M}$, if M does not lie entirely within vertical distance ε of ∂K_i^* or if its vertical projection does not lie entirely within B, then discard M from further consideration. For each remaining body M, select an arbitrary point from it, project this point vertically downward onto ∂K_i^*. Augment the point by associating it with any valid surface normal. Add the resulting augmented point to P_i. Repeating this process for all $M \in \mathcal{M}$ yields the desired set P_i.

To establish the correctness of this construction, consider any useful ε-cap C_p induced by an augmented point $p \in \partial K_i^*$ whose base Γ_p is of area at least t. Let h denote the supporting hyperplane at p, and let H denote the lower halfspace bounded by $h + \varepsilon$, which contains Γ_p. Recall that K_i^* is U-shaped and $O(1)$ steep. By Lemma 4.2 the vertical projection of C_p is bounded, and thus, by Lemma 3.1, $\mathrm{vol}(K_i^* \cap H) = \Omega(\varepsilon \cdot \mathrm{area}(\Gamma_p)) = \Omega(\varepsilon t)$. By choosing c suitably, we can ensure that $\mathrm{vol}(K_i^* \cap H) \geq v$. As observed in Lemma 3.2, it is possible to reduce the volume of this region by translating the halfspace downwards until the volume equals v. By Lemma 2.3(ii), there exists $M \in \mathcal{M}$ such that M is contained within the reduced region, and so $M \subseteq K_i^* \cap H$. (Note that M could not have been discarded in the construction process because, by Lemma 4.2, any useful ε-cap is contained within ball B.) It is easy to see that the point of P_i associated with M stabs C_p, as desired.

In order to bound the number of points in P_i, consider the set of points of K_i^* that lie within vertical distance ε of its boundary and whose vertical projection lies in B. The volume of this region is clearly $O(\varepsilon)$. By Lemma 2.3(i), each body M is of volume $\Omega(v) = \Omega(\varepsilon t)$. These bodies are pairwise disjoint, and (after discarding) each of them lies within the considered region. Therefore, by a simple packing argument, the number of such bodies is $O(\varepsilon/v) = O(1/t)$, as desired. $\square$

5 Putting the Pieces Together

Lemma 4.3 provides a set of augmented points $P_i \subset \partial K_i^*$ that stab all useful caps of K_i^* for which $\mathrm{area}(\Gamma_p)$ exceeds a given parameter t. Let's consider how to exploit this for stabbing dual caps in K_i. For each $p' \in P_i$, augment it by associating it with any valid surface normal, and let q' be the corresponding augmented point on ∂K_i. By basic properties of the dual transform, p' stabs a useful ε-cap $C_p(K_i^*)$ if and only if q' stabs the corresponding useful ε-dual cap $D_q(K_i)$. We can therefore map P_i to a set $Q_i' \subset \partial K_i$ in order to stab all the useful ε-dual caps of K_i that correspond to the ε-caps of K_i^* that are stabbed by P_i.

What area properties do these dual caps satisfy? Let $D_q(K_i)$ be any useful ε-dual cap and $C_p(K_i^*)$ be the corresponding useful cap. Let Δ_q' and Γ_p' denote the respective projections of the bases of $D_q(K_i)$ and $C_p(K_i^*)$. We claim that if $\text{area}(\Delta'_q) < c\varepsilon^{d-1}/t$, where c is a suitable constant, then D_q will be stabbed by some point in Q'_i. We will show that $\text{area}(\Gamma'_p) > t$, which will imply this claim since $\text{area}(\Gamma_p)$ is clearly more than $\text{area}(\Gamma'_p)$.

By Lemma 4.1, $\Delta'_q = \text{polar}_r(\Gamma'_p)$, for $r = \sqrt{\varepsilon}$. Recall that $\text{polar}_r(K)$ is a scaled copy of $\text{polar}_1(K)$ by $r^2 = \varepsilon$. Since these are $(d-1)$-dimensional sets, we have

$$\text{area}(\Delta'_q) \;=\; \text{area}(\varepsilon \cdot \text{polar}_1(\Gamma'_p)) \;=\; \varepsilon^{d-1} \cdot \text{area}(\text{polar}_1(\Gamma'_p)).$$

Thus, in order to bound $\text{area}(\Gamma'_p)$ in terms of $\text{area}(\Delta'_q)$, we need to establish a relationship between $\text{area}(\Gamma'_p)$ and $\text{area}(\text{polar}_1(\Gamma'_p))$. To do this, we make use of results on the Mahler volume. Recall that the Mahler volume of a convex body K is $\text{vol}(K) \cdot \text{vol}(\text{polar}_1(K))$. It is known that the Mahler volume is bounded below by a constant [24]. Applying this to the $(d-1)$-dimensional set Γ'_p, we have $\text{area}(\Gamma'_p) = \Omega(1/\text{area}(\text{polar}_1(\Gamma'_p)))$. Therefore, $\text{area}(\Gamma'_p) = \Omega(\varepsilon^{d-1}/\text{area}(\Delta'_q))$. Since $\text{area}(\Delta'_q) < c\varepsilon^{d-1}/t$, for a suitable c, it follows that $\text{area}(\Gamma'_p) > t$. Thus, adjusting for the constant factors, we have established the following analog to Lemma 3.2, but for small dual caps.

Lemma 5.1. *Let K be a convex body in $\mathbb{R}^d$, and consider any $\varepsilon > 0$ and $i \in I^{\pm}$. Given any $t > 0$, there exists a set of $O(1/t)$ augmented points $Q'_i \subset \partial K_i$ such that, for any useful ε-dual cap D_q induced by an augmented point $q \in \partial K_i^{(\varepsilon)}$, if $\text{area}(\Delta_q) \le \varepsilon^{d-1}/t$, there exists a point in Q'_i that stabs D_q.*

We are now ready to provide the proof of Theorem 1.1. Recall that $K_i = K_i^{(2\varepsilon)}$. By Lemma 2.2, it suffices to stab all the useful ε-dual caps of each of the K_i's, for $i \in I^{\pm}$. Fix any $i \in I^{\pm}$, and let $t = \sqrt{\text{area}(K)} \cdot \varepsilon^{(d-1)/2}$. For any augmented point $q \in \partial K_i^{(\varepsilon)}$, we say that the associated ε-dual cap is *large* if $\text{area}(\Delta_q) \ge t$, and otherwise it is *small*. We use two different strategies for stabbing the two types of dual caps.

For large dual caps, we apply Lemma 3.2 with the value of t defined above. This yields a set $Q_i \subset \partial K_i$ of size $O(\text{area}(K)/t) = O\big(\sqrt{\text{area}(K)}/\varepsilon^{(d-1)/2}\big)$, such that every useful large ε-dual cap is stabbed by one of these points.

In order to handle small dual caps, we apply Lemma 5.1 with the value of t set to $t' = c'\varepsilon^{(d-1)/2}/\sqrt{\text{area}(K)}$, for a suitably chosen constant c'. This yields a set $Q'_i \subset \partial K_i$ of size $O(1/t') = O\big(\sqrt{\text{area}(K)}/\varepsilon^{(d-1)/2}\big)$, such that every useful ε-dual cap whose base is of area at most $\varepsilon^{d-1}/t' = O\big(\sqrt{\text{area}(K)} \cdot \varepsilon^{(d-1)/2}\big)$ is stabbed. By choosing c' suitably, every useful small ε-dual cap is stabbed.

In summary, $Q_i \cup Q'_i$ is a set of size $O\big(\sqrt{\text{area}(K)}/\varepsilon^{(d-1)/2}\big)$ that stabs all useful ε-dual caps. By repeating this for all $i \in I^{\pm}$ and taking the union of all these sets, by Lemma 2.2, the resulting set provides the desired ε-approximation to K. This completes the proof of Theorem 1.1.

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] N. Alon and J. H. Spencer. *The probabilistic method*. Wiley, 2nd edition, 2008.

[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, T. Malamatos, and D. M. Mount. The effect of corners on the complexity of approximate range searching. *Discrete Comput. Geom.*, 41:398–443, 2009.

[5] S. Arya, D. M. Mount, and J. Xia. Tight lower bounds for halfspace range searching. *Discrete Comput. Geom.*, 2012. (doi: 10.1007/s00454-012-9412-x).

[6] I. Bárány. Intrinsic volumes and f-vectors of random polytopes. *Math. Ann.*, 285:671–699, 1989.

[7] I. Bárány. Random polytopes, convex bodies, and approximation. In W. Weil, editor, *Stochastic Geometry*, volume 1892 of *Lecture Notes in Mathematics*, pages 77–118. Springer, 2007.

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

[9] K. Böröczky, Jr. Approximation of general smooth convex bodies. *Advances in Mathematics*, 153:325–341, 2000.

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

[11] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite VC-dimension,. *Discrete Comput. Geom.*, 14:463–479, 1995.

[12] E. M. Bronshteyn and L. D. Ivanov. The approximation of convex sets by polyhedra. *Siberian Math. J.*, 16:852–853, 1976.

[13] E. M. Bronstein. Approximation of convex sets by polytopes. *Journal of Mathematical Sciences*, 153(6):727–762, 2008.

[14] B. Chazelle, L. J. Guibas, and D. T. Lee. The power of geometric duality. *BIT*, 25:76–90, 1985.

[15] K. L. Clarkson. Algorithms for polytope covering and approximation. In *Proc. 3th Workshop Algorithms Data Struct.*, pages 246–252, 1993.

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

[17] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 3rd edition, 2010.

[18] M. P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice Hall, 1976.

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

[20] H. G. Eggleston. *Convexity*. Cambridge University Press, 1958.

[21] 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.

[22] P. M. Gruber. Asymptotic estimates for best and stepwise approximation of convex bodies I. *Forum Math.*, 5:521–537, 1993.

[23] F. John. Extremum problems with inequalities as subsidiary conditions. In *Studies and Essays Presented to R. Courant on his 60th Birthday*, pages 187–204. Interscience Publishers, Inc., New York, 1948.

[24] G. Kuperberg. From the Mahler conjecture to Gauss linking integrals. *Geometric And Functional Analysis*, 18:870–892, 2008.

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

[26] D. E. McClure and R. A. Vitalie. Polygonal approximation of plane convex bodies. *J. Math. Anal. Appl.*, 51:326–358, 1975.

[27] J. Pach and G. Tardos. Tight lower bounds for the size of epsilon-nets. In *Proc. 27th Annu. Sympos. Comput. Geom.*, pages 458–463, 2011.

[28] R. Schneider. Polyhedral approximation of smooth convex bodies. *J. Math. Anal. Appl.*, 128:470–474, 1987.

[29] L. F. Toth. Approximation by polygons and polyhedra. *Bull. Amer. Math. Soc.*, 54:431–438, 1948.